

A CHARACTERISATION OF FIRST-ORDER CONSTRAINT SATISFACTION PROBLEMS*

BENOIT LAROSE^a, CYNTHIA LOTEN^b, AND CLAUDE TARDIF^c

^a Department of Mathematics and Statistics, Concordia University, 1455 de Maisonneuve West, Montréal, Qc, Canada, H3G 1M8
e-mail address: larose@mathstat.concordia.ca

^b Department of Mathematics and Statistics, University College of the Fraser Valley, 33844 King Rd Abbotsford, BC Canada V2S 7M8
e-mail address: cindy.loten@shaw.ca

^c Department of Mathematics and Computer Science, Royal Military College of Canada, PO Box 17000 Station “Forces”, Kingston, Ontario, Canada, K7K 7B4
e-mail address: Claude.Tardif@rmc.ca

ABSTRACT. We describe simple algebraic and combinatorial characterisations of finite relational core structures admitting finitely many obstructions. As a consequence, we show that it is decidable to determine whether a constraint satisfaction problem is first-order definable: we show the general problem to be **NP**-complete, and give a polynomial-time algorithm in the case of cores. A slight modification of this algorithm provides, for first-order definable CSP's, a simple poly-time algorithm to produce a solution when one exists. As an application of our algebraic characterisation of first order CSP's, we describe a large family of **L**-complete CSP's.

1. INTRODUCTION

The Constraint Satisfaction Problem (CSP) consists of determining, given a finite set of variables with constraints on these, whether there exists an assignment of values to these variables that satisfies all the given constraints. The great flexibility of this framework has made the CSP the focus of a great deal of attention from researchers in various fields (see for instance the recent survey [CJ06]). In general the problem is **NP**-complete, but restricting the type of constraint relations involved may yield tractable problems. In fact, Schaefer [Sch78] and more recently [ABISV05] have completely classified the complexity of Boolean CSP's and from their work it follows that Boolean CSP's are either trivial,

2000 ACM Subject Classification: F.2.2; F.4.m.

Key words and phrases: Constraint Satisfaction Problems, First Order Logic, Tree Duality, Finite Duality.

* A short version of this paper appeared in the Proceedings of the 21st Symposium on Logic in Computer Science (LICS 2006).

^a Research partially supported by NSERC, FQRNT and CRM.

^c Research partially supported by NSERC and ARP.

first-order definable, or complete (under AC^0 reductions) for one of the following standard classes of problems: **L**, **NL**, **P**, $\oplus\mathbf{L}$ and **NP**. One of the outstanding problems in the field is the so-called *dichotomy conjecture* [FV93] that states that every CSP should be either in **P** or **NP**-complete.

In this paper we adopt the convenient point of view offered in [FV98] where CSP's are viewed as homomorphism problems with a fixed target. In other words, if \mathbb{A} is a finite relational structure, then \mathbb{A} -CSP consists of all structures that admit a homomorphism to \mathbb{A} . Viewed this way, it becomes natural to ask which CSP's can be described in various logics. For instance in [LT07], the result of Allender et al. mentioned earlier is given a descriptive complexity analog, whereby it is shown that Boolean CSP's that lie in the classes **L** and **NL** are precisely those whose complement is describable in symmetric and linear Datalog respectively. Arguably the simplest CSP's (other than trivial ones) are those whose members are describable by a first-order sentence. A very natural question in the vein of the dichotomy conjecture is then the following: can we determine (easily) from the constraint relations whether a given CSP is first-order definable? Related questions for Datalog and its restrictions remain open [Dal05], [FV98]. An important first step in this direction is Atserias' result [Ats05] proving that FO-definable CSP's are precisely those with *finite duality*, i.e. those target structures \mathbb{A} for which there exists a finite set \mathcal{F} of structures such that \mathbb{B} admits no homomorphism to \mathbb{A} precisely if some structure in \mathcal{F} admits a homomorphism to \mathbb{B} . This result was followed closely by the more general result for homomorphism-closed classes by Rossman [Ros05].

In this paper, we give several equivalent characterisations of FO-definable CSP's. We first give a characterisation with an algebraic flavour: core structures with an FO-definable CSP are characterised by the existence of special near-unanimity operations preserving their basic relations (Theorem 4.5). For general structures, we prove that the problem of determining if \mathbb{A} -CSP is first-order definable is **NP**-complete (Theorem 6.1); if the structure \mathbb{A} is a core, then in fact there exists a simple polynomial-time algorithm to determine this (Theorem 6.2). We shall also describe in this case a simple algorithm that produces a solution in polynomial-time (Theorem 7.2). Let \mathbb{A} be a core structure such that \mathbb{A} -CSP is first-order definable, and let \mathbb{B} be a structure with the same universe, such that the basic relations of \mathbb{B} are constraint relations "inferred" from those of \mathbb{A} , i.e. each is describable by a primitive positive formula with atomic formulas of the form $\bar{x} \in \theta$ with θ a basic relation of \mathbb{A} ; these inferred relations play a crucial role in the study of the complexity of CSP's (see e.g. [CJ06]). It is known that \mathbb{B} -CSP is logspace reducible to \mathbb{A} -CSP [Jea98], but in general it will not be first-order definable. As a simple application of our algebraic characterisation of first-order definable CSP's (Corollary 4.3 and Proposition 4.4) we describe precisely which \mathbb{B} -CSP are first-order definable; the others turn out to be **L**-complete, with their complement definable in symmetric Datalog [ELT07], a fragment of linear Datalog.

To illustrate briefly the above results, we outline the algorithms in the special case of digraphs. For two vertices u, v of a digraph H , we say that v *dominates* u if every outneighbour of u is also an outneighbour of v and every inneighbour of u is also an inneighbour of v . If there exists a sequence $H = H_0, H_1, \dots, H_n = R$ of digraphs such that H_i is obtained from H_{i-1} by removing a dominated vertex for $i = 1, \dots, n$, we say that H *dismantles* to R . More generally, R is a *retract* of H if there exists a homomorphism from H to R whose restriction to R is the identity. The square R^2 of a digraph R has vertex set R^2 where the arcs are the couples $((u_0, u_1), (v_0, v_1))$ such that (u_0, v_0) and (u_1, v_1) are arcs of R , and its diagonal Δ_{R^2} is the set of vertices of R^2 with both coordinates equal.

The main algorithm to determine whether H -CSP is first-order definable proceeds as follows: in H^2 , remove any dominated element outside the diagonal, if any. Repeat this procedure until no element can be removed. If the resulting set is the diagonal, then the problem is first-order definable. Assuming that H is a core, i.e. that it has no proper retract, then the converse also holds.




Figure 1: The digraphs T_3 , P_3 and C .

In the figure above, T_3 is the transitive tournament on three vertices. In T_3^2 , the two isolated vertices $(0, 2)$, $(2, 0)$ are dominated by all other vertices, the sources $(0, 1)$, $(1, 0)$ are dominated by $(0, 0)$ and the sinks $(1, 2)$, $(2, 1)$ are dominated by $(2, 2)$. Hence T_3^2 dismantles to $\Delta_{T_3^2}$, which shows that T_3 -CSP is first-order definable. In fact it is well known that a directed graph G admits a homomorphism to T_3 if and only if there is no homomorphism from the directed 3-path P_3 to G , and this condition is described by the first-order sentence $\neg \exists a \exists b \exists c \exists d (A(a, b) \wedge A(b, c) \wedge A(c, d))$, where $A(x, y)$ denotes the existence of an arc from x to y . P_3 -CSP is not first-order definable; indeed the path P_3 is a core and P_3^2 can only be dismantled down to $P_3^2 \setminus \{(a, d), (d, a)\}$. The square of C cannot be dismantled to its diagonal, but C admits T_3 as a retract, whence C -CSP is first-order definable. Also, it is easy to check that $C \times T_3$ dismantles to the “graph” $\{(x, \phi(x)) : x \in C\}$ of a homomorphism $\phi : C \rightarrow T_3$. In Section 7, we will see that such dismantlings of products can always be used to produce solutions of first-order definable constraint satisfaction problems.

2. PRELIMINARIES

For basic notation and terminology we follow mainly [DKV02] and [NT00]. A *vocabulary* is a finite set $\sigma = \{R_1, \dots, R_m\}$ of *relation symbols*, each with an *arity* r_i assigned to it.

A σ -structure is a relational structure $\mathbb{A} = \langle A; R_1(\mathbb{A}), \dots, R_m(\mathbb{A}) \rangle$ where A is a non-empty set called the *universe* of \mathbb{A} , and $R_i(\mathbb{A})$ is an r_i -ary relation on A for each i . We will use the same capital letter in blackboard bold and slanted typeface to denote a structure and its universe respectively. The elements of $R_i(\mathbb{A})$, $1 \leq i \leq m$ will be called *hyperedges* of \mathbb{A} . For σ -structures \mathbb{A} and \mathbb{B} , a *homomorphism* from \mathbb{A} to \mathbb{B} is a map $f : A \rightarrow B$ such that $f(R_i(\mathbb{A})) \subseteq R_i(\mathbb{B})$ for all $1 \leq i \leq m$, where for any relation $R \in \sigma$ of arity r we have

$$f(R) = \{(f(x_1), \dots, f(x_r)) : (x_1, \dots, x_r) \in R\}.$$

A σ -structure \mathbb{B} is a *substructure* of a σ -structure \mathbb{A} if $B \subseteq A$ and the identity map on B is a homomorphism from \mathbb{B} to \mathbb{A} . For a subset B of A , the *substructure \mathbb{B} of \mathbb{A} induced by B* is the σ -structure with universe B with relations $R_i(\mathbb{B}) = R_i(\mathbb{A}) \cap B^{r_i}$ for every i . A substructure \mathbb{B} of \mathbb{A} is called a *retract* of \mathbb{A} if there exists a homomorphism ρ from \mathbb{A} to \mathbb{B} whose restriction to B is the identity; the map ρ is then called a *retraction*. A structure \mathbb{A} is called a *core* if it has no retract other than itself. It is well known (see [NT00]) that every (finite) σ -structure has a core which is unique up to isomorphism.

Let \mathbb{A} be a σ -structure. We define the *incidence multigraph* $\text{Inc}(\mathbb{A})$ of \mathbb{A} as the bipartite multigraph with parts A and $\text{Block}(\mathbb{A})$ which consists of all pairs (R, r) such that $R \in \sigma$ and $r \in R(A)$, and with edges $e_{a,i,B}$ joining $a \in A$ to $B = (R, (x_1, \dots, x_r)) \in \text{Block}(\mathbb{A})$ when $x_i = a$. This allows us to import some basic concepts from graph theory: the *distance* $d_{\mathbb{A}}(a, b)$ between two elements a and b of A is defined as half their distance in $\text{Inc}(\mathbb{A})$, the *diameter* of \mathbb{A} is defined as half the diameter of $\text{Inc}(\mathbb{A})$, and the *girth* of \mathbb{A} is defined as half the shortest length of a cycle in $\text{Inc}(\mathbb{A})$. In particular, \mathbb{A} has girth 1 if and only if $\text{Inc}(\mathbb{A})$ has parallel edges, and infinite girth if and only if $\text{Inc}(\mathbb{A})$ is acyclic. Notice in particular that tuples with repeated entries (such as (a, a, b)) create parallel edges and hence cycles; this property is not captured in the Gaifman graph. We'll require a finer notion of tree below and this explains why we choose this variant of a (multi)graph associated to a relational structure rather than the Gaifman graph.

Although this presentation of the girth differs from that given in [FV98], the concept is the same and we can use the following Erdős-type result.

Lemma 2.1 ([FV98] Theorem 5). Let \mathbb{A} and \mathbb{B} be σ -structures such that there exist no homomorphism from \mathbb{A} to \mathbb{B} . Then for any positive integer n there exists a σ -structure \mathbb{A}_n of girth greater than n such that there exists a homomorphism from \mathbb{A}_n to \mathbb{A} but no homomorphism from \mathbb{A}_n to \mathbb{B} .

Note that a σ -structure of large girth must have large diameter unless it is acyclic.

A *loop* in a σ -structure \mathbb{A} is an element $a \in A$ such that $(a, \dots, a) \in R_i(\mathbb{A})$ for any i ; equivalently, $a \in A$ is a loop if and only if for every σ -structure \mathbb{B} the constant map $\mathbb{B} \rightarrow \mathbb{A}$ with value a is a homomorphism. In particular, the image of a loop under a homomorphism is itself a loop. For an integer n the *n-link* of type $\sigma = \{R_1, \dots, R_m\}$ is the σ -structure

$$\mathbb{L}_n = \langle \{0, 1, \dots, n\}; R_1(\mathbb{L}_n), \dots, R_m(\mathbb{L}_n) \rangle,$$

such that $R_i(\mathbb{L}_n) = \cup_{j=1}^n \{j-1, j\}^{r_i}$ for $i = 1, \dots, m$ (where r_i is the arity of the relation R_i). Note that every $i \in \{0, 1, \dots, n\}$ is a loop in \mathbb{L}_n . A *link* in an arbitrary σ -structure is a homomorphic image of \mathbb{L}_n for some n . The term “path” is more common than “link”, but we chose the latter to make it clear that these are not trees in the sense defined below.

2.1. Trees. A σ -structure \mathbb{T} is called a σ -tree (or *tree* for short) if $\text{Inc}(\mathbb{T})$ is a tree, i.e. it is acyclic and connected. We require the following technical results:

Lemma 2.2. For every σ -tree \mathbb{T} with n hyperedges, there is a sequence $\mathbb{T} = \mathbb{T}_n, \mathbb{T}_{n-1}, \dots, \mathbb{T}_1$ of subtrees of \mathbb{T} with the following properties: for each $j = 1, \dots, n-1$

- (1) \mathbb{T}_j has j hyperedges;
- (2) \mathbb{T}_j is a subtree of \mathbb{T}_{j+1} ;
- (3) if (x_1, \dots, x_r) is the hyperedge of \mathbb{T}_{j+1} which does not belong to \mathbb{T}_j then there exists a unique index i such that x_i is in the universe of \mathbb{T}_j .

Proof. Let u_0, u_1, \dots, u_k be a path of maximal length in $\text{Inc}(\mathbb{T})$, where $\mathbb{T} = (T, R_1, \dots, R_m)$ has more than one hyperedge. If $u_0 = (R, (x_1)) \in \text{Block}(\mathbb{T})$, (R has to be a 1-ary relation of σ), we obtain a new tree \mathbb{T}' from \mathbb{T} by removing x_1 from R . If $u_0 \in T$, then $u_1 = (R, (x_1, \dots, x_r)) \in \text{Block}(\mathbb{T})$, and we obtain a new tree \mathbb{T}' by removing (x_1, \dots, x_r) from R and $\{x_1, \dots, x_m\} \setminus \{u_2\}$ from T . Repeating this procedure, we eventually obtain the desired decomposition. \square

Lemma 2.3. Let $\sigma = \{R_1, \dots, R_m\}$ be a vocabulary. Then for any integer n the number of core σ -trees of diameter at most n is finite.

Proof. We will show that the number t_n of core *rooted* trees in which the distance to the root is at most n is finite. Let m be the number of relations in σ and let r be the maximum arity of a relation in σ . We have $t_0 \leq 2^m$, with equality only if $r = 1$. Now suppose that t_{n-1} is finite. For a rooted tree \mathbb{T} in which the distance to the root u is at most n , we can encode each hyperedge $(x_1, \dots, x_{r'})$ to which u belongs by the name of the relation $R_i(\mathbb{T})$ containing it (there are at most m choices), the index i such that $u = x_i$ (there are at most r choices) and the trees rooted at $x_j, j \neq i$ branching away from u (there are at most t_{n-1}^{r-1} choices). If \mathbb{T} is a core, no two hyperedges can have the same label and \mathbb{T} is determined by its set of labels of hyperedges containing u . Therefore $t_n \leq 2^{m \cdot r \cdot t_{n-1}^{r-1}}$. \square

2.2. Complete sets of obstructions. The σ -structure \mathbb{B} is an *obstruction* for the σ -structure \mathbb{A} if there is no homomorphism from \mathbb{B} to \mathbb{A} . A family \mathcal{F} of obstructions for \mathbb{A} is called a *complete set of obstructions* if for every σ -structure \mathbb{B} that does not admit a homomorphism to \mathbb{A} there exists some $\mathbb{C} \in \mathcal{F}$ which admits a homomorphism to \mathbb{B} . The structure \mathbb{A} is said to have *tree duality* if it admits a complete set of obstructions consisting of trees, and *finite duality* if it admits a finite complete set of obstructions. According to [NT00], for every finite family \mathcal{F} of σ -trees, there exists a σ -structure $\mathbb{A}_{\mathcal{F}}$ which admits \mathcal{F} as a complete set of obstructions; and conversely every σ -structure \mathbb{A} with finite duality admits a finite complete set of obstructions consisting of trees. Thus the structures with finite duality form a subclass of the structures with tree duality, and there is one such core structure for every finite set of tree obstructions.

An obstruction \mathbb{B} for \mathbb{A} is called *critical* if every proper substructure of \mathbb{B} admits a homomorphism to \mathbb{A} . It is clear that a critical obstruction is a core, and that every obstruction contains, as a substructure, a critical obstruction.

Lemma 2.4. A σ -structure \mathbb{A} has finite duality if and only if there is an upper bound on the diameter of its critical obstructions.

Proof. Clearly, if \mathbb{A} has finite duality, then the maximum diameter of an obstruction in a finite complete set of obstructions for \mathbb{A} is an upper bound on the diameter of all critical obstructions for \mathbb{A} . Conversely, suppose that the critical obstructions for \mathbb{A} have diameter at most m . Let \mathcal{F} be the set of core σ -trees of diameter at most m which do not admit a homomorphism to \mathbb{A} . By Lemma 2.3, \mathcal{F} is finite. By Lemma 2.1, for any σ -structure \mathbb{B} which does not admit a homomorphism to \mathbb{A} , there exists a structure \mathbb{C} of girth at least $2m + 2$ which admits a homomorphism to \mathbb{B} but not to \mathbb{A} . A critical obstruction for \mathbb{A} contained in \mathbb{C} cannot contain a cycle hence it must be a tree \mathbb{T} of diameter at most m . Therefore $\mathbb{T} \in \mathcal{F}$; this shows that \mathcal{F} is a finite complete set of obstructions for \mathbb{A} . \square

For a σ -structure \mathbb{A} , the problem \mathbb{A} -CSP consists of determining whether an input structure \mathbb{B} admits a homomorphism to \mathbb{A} . It is said to be *first-order definable* if there exists a first-order sentence Φ (in the language of σ) which is true on \mathbb{B} if and only if \mathbb{B} admits a homomorphism to \mathbb{A} . By a result of Atserias [Ats05], \mathbb{A} -CSP is first-order definable if and only if \mathbb{A} has finite duality, hence we have the following equivalences:

Theorem 2.5. Let \mathbb{A} be a σ -structure. Then the following are equivalent.

1. \mathbb{A} -CSP is first-order definable;
2. \mathbb{A} has finite duality;
3. \mathbb{A} has a finite complete set of obstructions consisting of trees;
4. The critical obstructions of \mathbb{A} have bounded diameter.

We are mostly interested in the “meta-problem” of deciding whether an input structure \mathbb{A} has a first-order definable CSP. The equivalences of Theorem 2.5 are not a usable decision procedure, but they will be used in the next two sections to find such a procedure. As a benchmark we state here Feder and Vardi’s decision procedure for tree duality. Given a structure $\mathbb{A} = \langle A; R_1(\mathbb{A}), \dots, R_m(\mathbb{A}) \rangle$, we define the structure $\mathcal{U}(\mathbb{A}) = \langle U; R_1(\mathcal{U}(\mathbb{A})), \dots, R_m(\mathcal{U}(\mathbb{A})) \rangle$, where U is the set of all nonempty subsets of A , and for $i = 1, \dots, m$, $R_i(\mathcal{U}(\mathbb{A}))$ is the set of all r_i -tuples (X_1, \dots, X_{r_i}) such that for all $j \in \{1, \dots, r_i\}$ and $x_j \in X_j$ there exist $x_k \in X_k, k \in \{1, \dots, r_i\} \setminus \{j\}$ such that $(x_1, \dots, x_{r_i}) \in R_i(\mathbb{A})$.

Theorem 2.6 ([FV98] Theorem 21). A σ -structure \mathbb{A} has tree duality if and only if there exists a homomorphism from $\mathcal{U}(\mathbb{A})$ to \mathbb{A} .

This proves that determining whether a given structure \mathbb{A} has tree duality is decidable, since a search for a homomorphism from $\mathcal{U}(\mathbb{A})$ to \mathbb{A} can be done in finite time. In section 4, we provide similar “construction-and-homomorphism” characterisations of first-order definable constraint satisfaction problems. However, we must first clear up a technical point concerning tree duality: indeed, Feder and Vardi’s definition of a tree given in [FV98] (at the bottom of page 79) is slightly more general than the one given here, as it allows parallel edges in the incidence multigraph of a tree. Nonetheless the corresponding concepts of “tree duality” turn out to be equivalent, as we show in the next section.

3. TREE DUALITY

In this section we prove that the notion of tree duality is the same whether we use the notion of tree as defined here or as defined in [FV98].

Lemma 3.1. Let \mathbb{T} be a tree. Then \mathbb{T} admits a homomorphism to \mathbb{A} if and only if \mathbb{T} admits a homomorphism to $\mathcal{U}(\mathbb{A})$.

Proof. First note that the singletons induce a copy of \mathbb{A} in $\mathcal{U}(\mathbb{A})$. Thus, if a tree \mathbb{T} admits a homomorphism to \mathbb{A} , it also admits a homomorphism to $\mathcal{U}(\mathbb{A})$.

Conversely, suppose that $\phi : \mathbb{T} \rightarrow \mathcal{U}(\mathbb{A})$ is a homomorphism. Let $\mathbb{T} = \mathbb{T}_n, \mathbb{T}_{n-1}, \dots, \mathbb{T}_1$ be a decomposition of \mathbb{T} as described in Lemma 2.2. We can define a sequence of homomorphisms $\psi_i : \mathbb{T}_i \rightarrow \mathbb{A}$ as follows: \mathbb{T}_1 has a single hyperedge $(x_1, \dots, x_r) \in R(\mathbb{A})$ for some $R \in \sigma$. Since ϕ is a homomorphism, by definition of $\mathcal{U}(\mathbb{A})$ there exist $\psi_1(x_j) \in \phi(x_j)$, $j = 1, \dots, r$, such that $(\psi_1(x_1), \dots, \psi_1(x_r)) \in R(\mathcal{U}(\mathbb{A}))$. This defines a homomorphism $\psi_1 : \mathbb{T}_1 \rightarrow \mathbb{A}$.

Now suppose that $i < n$ and $\psi_i : \mathbb{T}_i \rightarrow \mathbb{A}$ is already defined. \mathbb{T}_{i+1} is obtained from \mathbb{T}_i by adding one hyperedge $(x_1, \dots, x_r) \in R(\mathbb{T})$ for some $R \in \sigma$, where for exactly one index j_0 , x_{j_0} belongs to the universe of \mathbb{T}_i . Since ϕ is a homomorphism, by definition of $\mathcal{U}(\mathbb{A})$ there exist $\psi'_{i+1}(x_j) \in \phi(x_j)$, $j = 1, \dots, r$, such that $(\psi'_{i+1}(x_1), \dots, \psi'_{i+1}(x_r)) \in R(\mathcal{U}(\mathbb{A}))$ and $\psi'_{i+1}(x_{j_0}) = \psi_i(x_{j_0})$. Then $\psi_i \cup \psi'_{i+1} = \psi_{i+1}$ is a well defined homomorphism from \mathbb{T}_{i+1} to \mathbb{A} . Continuing in this way, we eventually define a homomorphism $\psi = \psi_n$ from \mathbb{T} to \mathbb{A} . \square

The \mathbb{A} hyperedge consistency check is the following polynomial-time algorithm. At the start every element $b \in B$ is assigned a list consisting of all the “plausible” images of b under such a homomorphism; initially this list is A . Then, the elements of B are cyclically inspected to check whether their lists are still consistent with the local information: For every element b , an element a in the current list of b is removed if there exists a relation $R \in \sigma$ and $(b_1, \dots, b_r) \in R(\mathbb{B})$ with $b_j = b$ for some j such that there exists no $(a_1, \dots, a_r) \in R(\mathbb{A})$ with $a_j = a$ and a_i in the list of b_i for $i \neq j$. The process continues until the lists stabilise. If at some point the list of an element becomes empty, the \mathbb{A} hyperedge consistency check is said to fail on \mathbb{B} , and otherwise it is said to succeed on \mathbb{B} . The following result gives an interpretation of these possible outcomes.

Lemma 3.2. *The \mathbb{A} hyperedge consistency check succeeds on \mathbb{B} if and only if there exists a homomorphism from \mathbb{B} to $\mathcal{U}(\mathbb{A})$, and it fails on \mathbb{B} if and only if there exists a tree \mathbb{T} which admits a homomorphism to \mathbb{B} but no homomorphism to \mathbb{A} .*

Proof. If the \mathbb{A} hyperedge consistency check succeeds on \mathbb{B} , then by definition of $\mathcal{U}(\mathbb{A})$ the map $\phi : \mathbb{B} \rightarrow \mathcal{U}(\mathbb{A})$ assigning to every $b \in B$ its final list $\phi(b)$ is a homomorphism. Conversely, if $\phi : \mathbb{B} \rightarrow \mathcal{U}(\mathbb{A})$ is a homomorphism, then again by definition of $\mathcal{U}(\mathbb{A})$, the \mathbb{A} hyperedge consistency check will never eliminate an element $c \in \phi(b)$ from the list of any element b of \mathbb{B} , hence it will succeed.

If there exists a tree \mathbb{T} which admits a homomorphism to \mathbb{B} but not to \mathbb{A} , then by Lemma 3.1, \mathbb{T} does not admit a homomorphism to $\mathcal{U}(\mathbb{A})$, thus \mathbb{B} does not admit a homomorphism to $\mathcal{U}(\mathbb{A})$. By the previous paragraph this implies that the \mathbb{A} hyperedge consistency check must fail on \mathbb{B} . Conversely, suppose that the \mathbb{A} hyperedge consistency check fails on \mathbb{B} . We construct a tree \mathbb{T} as follows while running the hyperedge consistency check.

When deleting an element a from the list of an element b of \mathbb{B} , we define a rooted tree $\mathbb{T}_{a,b}$ with root $r_{a,b}$ with the following properties.

- (i) There is a homomorphism from $\mathbb{T}_{a,b}$ to \mathbb{B} mapping $r_{a,b}$ to b ,
- (ii) there is no homomorphism from $\mathbb{T}_{a,b}$ to \mathbb{A} mapping $r_{a,b}$ to a .

Indeed, a is deleted from the list of b because we found a relation $R \in \sigma$ and $(b_1, \dots, b_r) \in R(\mathbb{B})$ with $b_j = b$ for some j such that there exists no $(a_1, \dots, a_r) \in R(\mathbb{A})$ with $a_j = a$ and a_i in the list of b_i for $i \neq j$. We then put elements c_1, \dots, c_r in the universe of $\mathbb{T}_{a,b}$ and $(c_1, \dots, c_r) \in R(\mathbb{T}_{a,b})$, and select the root $r_{a,b} = c_j$. If $R(\mathbb{A})$ does not contain any r -tuple

(a_1, \dots, a_r) such that $a_j = a$, then we are done, since $\mathbb{T}_{a,b}$ clearly satisfies properties (i) and (ii). Otherwise, for every (a_1, \dots, a_r) such that $a_j = a$, there exists at least one index i such that a_i is already removed from the list of b_i whence the tree \mathbb{T}_{a_i,b_i} with properties (i) and (ii) is already defined; we then add a copy of \mathbb{T}_{a_i,b_i} to $\mathbb{T}_{a,b}$, by identifying r_{a_i,b_i} to c_i . Thus we get a tree $\mathbb{T}_{a,b}$ such that the map $\phi : \{c_1, \dots, c_r\} \rightarrow \{b_1, \dots, b_r\}$ defined by $\phi(c_i) = b_i$ extends to a homomorphism from $\mathbb{T}_{a,b}$ to \mathbb{B} . However, any homomorphism from $\mathbb{T}_{a,b}$ to \mathbb{A} mapping $r_{a,b}$ to a would also map the root of some previously defined \mathbb{T}_{a_i,b_i} to a_i , which is impossible. Thus $\mathbb{T}_{a,b}$ satisfies properties (i) and (ii).

When the list of some element b of \mathbb{B} becomes empty, we can construct a tree \mathbb{T} by identifying the roots of all the trees $\mathbb{T}_{a,b}$ to a new element r . We then find a homomorphism from \mathbb{T} to \mathbb{B} by mapping r to b and extending independently on each $\mathbb{T}_{a,b}$. However a homomorphism from \mathbb{T} to \mathbb{A} would need to map r to some element a , hence induce a homomorphism from $\mathbb{T}_{a,b}$ to \mathbb{A} mapping its root to a , which is impossible. Therefore, when the hyperedge consistency check fails, there exists a tree \mathbb{T} which admits a homomorphism to \mathbb{B} but not to \mathbb{A} . \square

In terms of dualities, these results can be summarized as follows.

Theorem 3.3. *For a σ -structure \mathbb{A} , the following properties are equivalent:*

- (i) \mathbb{A} has tree duality,
- (ii) $\mathcal{U}(\mathbb{A})$ admits a homomorphism to \mathbb{A} ,
- (iii) The \mathbb{A} hyperedge consistency check decides the \mathbb{A} -CSP problem.

Proof.

- (i) \Rightarrow (ii) Suppose that \mathbb{A} has tree duality. By Lemma 3.1, there does not exist a tree that admits a homomorphism to $\mathcal{U}(\mathbb{A})$ but not to \mathbb{A} , hence there exists a homomorphism from $\mathcal{U}(\mathbb{A})$ to \mathbb{A} .
- (ii) \Rightarrow (iii) By Lemma 3.2 the \mathbb{A} hyperedge consistency check decides the $\mathcal{U}(\mathbb{A})$ -CSP problem. If there exists a homomorphism from $\mathcal{U}(\mathbb{A})$ to \mathbb{A} , then since there also exists a homomorphism from \mathbb{A} to $\mathcal{U}(\mathbb{A})$ the \mathbb{A} -CSP problem is equivalent to the $\mathcal{U}(\mathbb{A})$ -CSP problem.
- (iii) \Rightarrow (i) By Lemma 3.2, when the \mathbb{A} hyperedge consistency check fails on a structure \mathbb{B} , there exists a tree \mathbb{T} which admits a homomorphism to \mathbb{B} but none to \mathbb{A} . Thus if the \mathbb{A} hyperedge consistency check decides the \mathbb{A} -CSP problem, then \mathbb{A} has tree duality. \square

Property (ii) shows that these three properties are decidable, since a greedy search for a homomorphism from $\mathcal{U}(\mathbb{A})$ to \mathbb{A} can be done in finite time. Property (iii) gives a polynomial algorithm for the corresponding CSP's, and property (i) shows that these problems contain the class of CSP's with finite duality, that is, the first-order decidable CSP's.

4. CONSTRUCTIONS

4.1. Quotients. Let $\mathbb{A} = \langle A; R_1(\mathbb{A}), \dots, R_m(\mathbb{A}) \rangle$ be a σ -structure and \sim an equivalence relation on A . For $a \in A$ we denote a/\sim the \sim -equivalence class containing a . The *quotient* \mathbb{A}/\sim of \mathbb{A} under \sim is the σ -structure whose universe is the set of \sim -equivalence classes, where for $i = 1, \dots, m$ we have $(C_1, \dots, C_{r_i}) \in R_i(\mathbb{A}/\sim)$ if and only if there exist $a_j \in C_j$,

$j = 1, \dots, r_i$ such that $(a_1, \dots, a_{r_i}) \in R_i(\mathbb{A})$. Note that the quotient map $q : \mathbb{A} \rightarrow \mathbb{A}/\sim$ where $q(a) = a/\sim$ is a homomorphism; in fact for every homomorphism $\phi : \mathbb{A} \rightarrow \mathbb{B}$, there is a natural equivalence \sim (the “kernel” of ϕ) on A and an injective homomorphism $\psi : \mathbb{A}/\sim \rightarrow \mathbb{B}$ such that $\phi = \psi \circ q$.

Here we give a first application of quotients to reveal an important structural property of cores with tree duality. A σ -structure \mathbb{A} is called *rigid* if the identity is the only homomorphism from \mathbb{A} to itself.

Lemma 4.1. Let \mathbb{A} be a core with tree duality. Then \mathbb{A} is rigid.

Proof. Suppose that $\tau : \mathbb{A} \rightarrow \mathbb{A}$ is a homomorphism. Since \mathbb{A} is a core, τ is an automorphism of \mathbb{A} hence we can define an equivalence relation \sim on A by putting $a \sim b$ if there exists an integer p such that $\tau^p(a) = b$. We will show that every tree which admits a homomorphism to \mathbb{A}/\sim also admits a homomorphism to \mathbb{A} .

Let \mathbb{T} be a tree which admits a homomorphism $\psi : \mathbb{T} \rightarrow \mathbb{A}/\sim$. Let $\mathbb{T} = \mathbb{T}_n, \mathbb{T}_{n-1}, \dots, \mathbb{T}_1$ be the sequence of Lemma 2.2. For $k = 1, \dots, n$, the restriction of ψ to the universe of \mathbb{T}_k is a homomorphism $\psi_k : \mathbb{T}_k \rightarrow \mathbb{A}/\sim$; we recursively define a sequence $\phi_k : \mathbb{T}_k \rightarrow \mathbb{A}$ of homomorphisms such that $\psi_k = q \circ \phi_k$, where q is the quotient map from \mathbb{A} to \mathbb{A}/\sim . First, \mathbb{T}_1 has just one hyperedge $(x_1, \dots, x_{r_i}) \in R_i(\mathbb{T})$ for some i , and $(\psi(x_1), \dots, \psi(x_{r_i})) \in R_i(\mathbb{A}/\sim)$. By definition of quotients this means that there exist $y_j \in \psi(x_j), j = 1, \dots, r_i$ such that $(y_1, \dots, y_{r_i}) \in R_i(\mathbb{A})$, thus we can define $\phi_1 : \mathbb{T}_1 \rightarrow \mathbb{A}$ by $\phi_1(x_j) = y_j$. Now suppose that $\phi_{k-1} : \mathbb{T}_{k-1} \rightarrow \mathbb{A}/\sim$ is already defined. \mathbb{T}_k is obtained from \mathbb{T}_{k-1} by adding an hyperedge $(x_1, \dots, x_{r_i}) \in R_i(\mathbb{T})$ which has only one coordinate x_ℓ in the universe of \mathbb{T}_{k-1} . Again we have $(\psi(x_1), \dots, \psi(x_{r_i})) \in R_i(\mathbb{A}/\sim)$ and there exist $y_j \in \psi(x_j), j = 1, \dots, r_i$ such that $(y_1, \dots, y_{r_i}) \in R_i(\mathbb{A})$. Put $a = \phi_{k-1}(x_\ell) \in \mathbb{A}/\sim = \psi(x_\ell)$. Then $y_\ell \sim a$ hence by the definition of \sim there exists a power p such that $\tau^p(y_\ell) = a$. Since τ is a homomorphism, we then have $(\tau^p(y_1), \dots, \tau^p(y_{r_i})) \in R_i(\mathbb{A})$, and we can extend the definition of ϕ_{k-1} to that of $\phi_k : \mathbb{T}_k \rightarrow \mathbb{A}$ by putting $\phi_k(z) = \phi_{k-1}(z)$ if z is in the universe of \mathbb{T}_{k-1} , and $\phi_k(x_j) = \tau^p(y_j), j = 1, \dots, r_i$. Indeed ϕ_k is well defined since both definitions coincide on x_ℓ , it is a homomorphism since it preserves $(x_1, \dots, x_{r_i}) \in R_i(\mathbb{T})$ in addition to all the hyperedges preserved by ϕ_{k-1} , and $\phi_k(z) \in \psi_k(z)$ for all z in the universe of \mathbb{T}_k whence $\psi_k(z) = q \circ \phi_k(z)$. In this way we eventually define a homomorphism $\phi = \phi_n$ from $\mathbb{T} = \mathbb{T}_n$ to \mathbb{A} .

Hence every tree which admits a homomorphism to \mathbb{A}/\sim also admits a homomorphism to \mathbb{A} . Since \mathbb{A} has tree duality this implies that \mathbb{A}/\sim admits a homomorphism to \mathbb{A} . Since \mathbb{A} is a core which admits a homomorphism to \mathbb{A}/\sim , this implies that \sim cannot identify vertices, whence τ is the identity. \square

4.2. Products and powers. Given two σ -structures $\mathbb{A} = \langle A; R_1(\mathbb{A}), \dots, R_m(\mathbb{A}) \rangle$ and $\mathbb{B} = \langle B; R_1(\mathbb{B}), \dots, R_m(\mathbb{B}) \rangle$ their *product* is the σ -structure

$$\mathbb{A} \times \mathbb{B} = \langle A \times B; R_1(\mathbb{A} \times \mathbb{B}), \dots, R_m(\mathbb{A} \times \mathbb{B}) \rangle,$$

where for $i = 1, \dots, m$, $R_i(\mathbb{A} \times \mathbb{B})$ consists of all tuples $((a_1, b_1), \dots, (a_{r_i}, b_{r_i}))$ such that $(a_1, \dots, a_{r_i}) \in R_i(\mathbb{A})$ and $(b_1, \dots, b_{r_i}) \in R_i(\mathbb{B})$. Both projections $\pi_1 : \mathbb{A} \times \mathbb{B} \rightarrow \mathbb{A}$ and $\pi_2 : \mathbb{A} \times \mathbb{B} \rightarrow \mathbb{B}$ are homomorphism and in general for any σ -structure \mathbb{C} and any pair $\phi_1 : \mathbb{C} \rightarrow \mathbb{A}, \phi_2 : \mathbb{C} \rightarrow \mathbb{B}$ of homomorphisms there is a unique homomorphism $\phi : \mathbb{C} \rightarrow \mathbb{A} \times \mathbb{B}$ such that $\phi_1 = \pi_1 \circ \phi$ and $\phi_2 = \pi_2 \circ \phi$. The product is associative; the n -th power \mathbb{A}^n of \mathbb{A}

is the product of n copies of \mathbb{A} . For any $n \geq 1$ an n -ary operation on \mathbb{A} is a homomorphism from \mathbb{A}^n to \mathbb{A} .

The *one-tolerant n -th power* ${}^1\mathbb{A}^n$ of \mathbb{A} is the σ -structure $\langle A^n; R_1({}^1\mathbb{A}^n), \dots, R_m({}^1\mathbb{A}^n) \rangle$ where for $i = 1, \dots, m$, $R_i({}^1\mathbb{A}^n)$ consists of tuples $((a_{1,1}, \dots, a_{1,n}), \dots, (a_{r_i,1}, \dots, a_{r_i,n}))$ such that $|\{k : (a_{1,k}, \dots, a_{r_i,k}) \in R_i(\mathbb{A})\}| \geq n - 1$. In other words, ${}^1\mathbb{A}^n$ is obtained from \mathbb{A}^n by adding to $R_i(\mathbb{A}^n)$ all hyperedges that are mapped to $R_i(\mathbb{A})$ by at least $n - 1$ of the projections. In particular, the projections are not homomorphisms from ${}^1\mathbb{A}^n$ to \mathbb{A} hence ${}^1\mathbb{A}^n$ does not necessarily admit a homomorphism to \mathbb{A} . However notice that removal of a coordinate is a homomorphism from ${}^1\mathbb{A}^{n+1}$ to ${}^1\mathbb{A}^n$.

Lemma 4.2. There exists a homomorphism from ${}^1\mathbb{A}^{n+1}$ to \mathbb{A} if and only if the critical obstructions of \mathbb{A} have at most n hyperedges.

Proof. Let \mathbb{C} be a critical obstruction of \mathbb{A} with m distinct hyperedges e_1, \dots, e_m , $m > n$. Then for $k = 1, \dots, m$, the σ -structure \mathbb{C}_k obtained from \mathbb{C} by removing e_k (without changing the universe) admits a homomorphism ϕ_k to \mathbb{A} . By definition of ${}^1\mathbb{A}^m$, the map $\phi = (\phi_1, \dots, \phi_m)$ is a homomorphism from \mathbb{C} to ${}^1\mathbb{A}^m$. Therefore there is no homomorphism from ${}^1\mathbb{A}^m$ to \mathbb{A} , and in particular none from ${}^1\mathbb{A}^{n+1}$ to \mathbb{A} .

Conversely, suppose that there is no homomorphism from ${}^1\mathbb{A}^{n+1}$ to \mathbb{A} . Then there exists a critical obstruction \mathbb{C} of \mathbb{A} which admits a homomorphism ϕ to ${}^1\mathbb{A}^{n+1}$. For every coordinate $k = 1, \dots, n + 1$, there exists an hyperedge e_k of \mathbb{C} which is not respected by $\pi_k \circ \phi$, since $\pi_k \circ \phi$ is not a homomorphism from \mathbb{C} to \mathbb{A} . By the definition of ${}^1\mathbb{A}^{n+1}$, e_k is respected by $\pi_j \circ \phi$ for every $j \neq k$, whence $e_j \neq e_k$ for $j \neq k$. Therefore \mathbb{C} has at least $n + 1$ hyperedges. \square

Corollary 4.3. A σ -structure \mathbb{A} has finite duality if and only if there exists a positive integer n such that ${}^1\mathbb{A}^n$ admits a homomorphism to \mathbb{A} .

Note that the homomorphisms from 1-tolerant powers of \mathbb{A} to \mathbb{A} are operations on \mathbb{A} . For $n \geq 3$, an operation $\phi : \mathbb{A}^n \rightarrow \mathbb{A}$ is called a *near unanimity operation* if it satisfies the identities

$$\phi(y, x, x, \dots, x) = \phi(x, y, x, \dots, x) = \dots = \phi(x, x, x, \dots, y) = x.$$

Lemma 4.4. Let \mathbb{A} be a core with finite duality. Then every homomorphism from a 1-tolerant power of \mathbb{A} to \mathbb{A} is a near unanimity operation.

Proof. Let $\phi : {}^1\mathbb{A}^n \rightarrow \mathbb{A}$ be a homomorphism. For every $y \in A$ and $k \in \{1, \dots, n\}$, consider the homomorphism $\psi_{y,k} : \mathbb{A} \rightarrow {}^1\mathbb{A}^n$ defined by $\psi_{y,k}(x) = (x_1, \dots, x_n)$ where $x_j = y$ if $j = k$ and $x_j = x$ otherwise. By Lemma 4.1, \mathbb{A} is rigid whence the map $\phi \circ \psi_{y,k} : \mathbb{A} \rightarrow \mathbb{A}$ is the identity. Thus for every $x, y \in A$ and $k \in \{1, \dots, n\}$ we have $\phi(\psi_{y,k}(x)) = x$, and this is precisely the definition of a near unanimity operation. \square

We say that a structure \mathbb{A} *admits* an operation $f : A^n \rightarrow A$, or equivalently that f *preserves* the basic relations of \mathbb{A} if f is a homomorphism from \mathbb{A}^n to \mathbb{A} .

Corollary 4.5. Every core relational structure with a first-order definable CSP admits a near unanimity operation.

4.3. Products of links and squares. Recall from Section 2 that the n -link \mathbb{L}_n of type σ has universe $\{0, 1, \dots, n\}$. For a σ -structure \mathbb{C} , a map ϕ from its universe to $\{0, 1, \dots, n\}$ is a homomorphism from \mathbb{C} to \mathbb{L}_n if and only if $|\phi(x) - \phi(y)| \leq 1$ whenever x and y are in a common hyperedge.

Given a σ -structure $\mathbb{A} = \langle A; R_1(\mathbb{A}), \dots, R_m(\mathbb{A}) \rangle$, note that the product $\mathbb{L}_n \times \mathbb{A}^2$ has diameter at least n since for any $a, a', b, b' \in A$ the distance between $(0, a, b)$ and (n, a', b') is at least n . (The distance could even be infinite, that is, $(0, a, b)$ and (n, a', b') could lie in different connected components.) Let \sim_n be the equivalence relation defined on $\mathbb{L}_n \times \mathbb{A}^2$ by

$$(k, a, b) \sim_n (k', a', b') \equiv \begin{cases} (k, a, b) = (k', a', b') \\ \text{or } k = k' = 0 \text{ and } a = a' \\ \text{or } k = k' = n \text{ and } b = b'. \end{cases}$$

Note that $\mathbb{L}_n \times \mathbb{A}^2 / \sim_n$ also has diameter at least n .

Lemma 4.6. The substructures \mathbb{B}_0 and \mathbb{B}_n of $\mathbb{L}_n \times \mathbb{A}^2 / \sim_n$ induced by $B_0 = \{(k, a, b) / \sim_n : k \neq 0\}$ and $B_n = \{(k, a, b) / \sim_n : k \neq n\}$ respectively both admit homomorphisms to \mathbb{A} .

Proof. On B_0 we can define a map ϕ to A by $\phi((k, a, b) / \sim_n) = b$. We show that ϕ is a homomorphism from \mathbb{B}_0 to \mathbb{A} . For $R_i \in \sigma$ and $((k_1, a_1, b_1) / \sim_n, \dots, (k_{r_i}, a_{r_i}, b_{r_i}) / \sim_n) \in R_i(\mathbb{B}_0)$, there exist $(k'_j, a'_j, b'_j) \in (k_j, a_j, b_j) / \sim_n$, $j = 1, \dots, r_i$, with $((k'_1, a'_1, b'_1), \dots, (k'_{r_i}, a'_{r_i}, b'_{r_i})) \in R_i(\mathbb{L}_n \times \mathbb{A}^2)$. We then have that

$$(\phi((k_1, a_1, b_1) / \sim_n), \dots, \phi((k_{r_i}, a_{r_i}, b_{r_i}) / \sim_n))$$

is equal to (b'_1, \dots, b'_{r_i}) which is in $R_i(\mathbb{A})$, thus ϕ is a homomorphism. Similarly, we can define a homomorphism $\psi : \mathbb{B}_n \rightarrow \mathbb{A}$ by $\psi((k, a, b) / \sim_n) = a$. \square

Proposition 4.1. A σ -structure \mathbb{A} has critical obstructions of bounded diameter if and only if there exists a positive integer n such that $\mathbb{L}_n \times \mathbb{A}^2 / \sim_n$ admits a homomorphism to \mathbb{A} .

Proof. By the previous lemma, any critical obstruction of \mathbb{A} contained in $\mathbb{L}_n \times \mathbb{A}^2 / \sim_n$ must contain an element with first coordinate 0 and an element with first coordinate n (the first coordinates are invariants of \sim_n -equivalence classes) thus have diameter at least n . Hence if n is larger than the diameter of all the critical obstructions of \mathbb{A} , then $\mathbb{L}_n \times \mathbb{A}^2 / \sim_n$ admits a homomorphism to \mathbb{A} .

Now suppose that \mathbb{A} has critical obstructions of arbitrarily large diameter. We will show that for every integer n there exists an obstruction \mathbb{C} of \mathbb{A} which admits a homomorphism to $\mathbb{L}_n \times \mathbb{A}^2 / \sim_n$. Let $\mathbb{C} = \langle C; R_1(\mathbb{C}), \dots, R_m(\mathbb{C}) \rangle$ be an obstruction of \mathbb{A} with diameter at least $n + 2$. Let x and y be elements of C at distance $n + 2$, and $\mathbb{C}_x, \mathbb{C}_y$ the substructures of \mathbb{C} induced respectively by $C \setminus \{x\}$ and $C \setminus \{y\}$. Fix homomorphisms $\alpha : \mathbb{C}_y \rightarrow \mathbb{A}$ and $\beta : \mathbb{C}_x \rightarrow \mathbb{A}$ and define $\kappa : C \rightarrow \{0, \dots, n\}$ by

$$\kappa(z) = \begin{cases} 0 & \text{if } z = x, \\ d_{\mathbb{C}}(x, z) - 1 & \text{if } d_{\mathbb{C}}(x, z) \leq n + 1 \text{ and } z \neq x, \\ n & \text{if } d_{\mathbb{C}}(x, z) \geq n + 2; \end{cases}$$

note that κ is a homomorphism from \mathbb{C} to \mathbb{L}_n . We fix an element $p \in A$ and define a map ϕ from C to the universe of $\mathbb{L}_n \times \mathbb{A}^2 / \sim_n$ by

$$\phi(z) = \begin{cases} (\kappa(z), \alpha(z), \beta(z)) / \sim_n & \text{if } z \neq x, y, \\ (\kappa(z), \alpha(z), p) / \sim_n & \text{if } z = x, \\ (\kappa(z), p, \beta(z)) / \sim_n & \text{if } z = y. \end{cases}$$

We will show that ϕ is a homomorphism from \mathbb{C} to $\mathbb{L}_n \times \mathbb{A}^2 / \sim_n$.

Let (z_1, \dots, z_{r_i}) be in $R_i(\mathbb{C})$ for some $R_i \in \sigma$. If $z_j \notin \{x, y\}$ for all $j \in \{1, \dots, r_i\}$, then $(\phi(z_1), \dots, \phi(z_{r_i})) = ((\kappa(z_1), \alpha(z_1), \beta(z_1)) / \sim_n, \dots, (\kappa(z_{r_i}), \alpha(z_{r_i}), \beta(z_{r_i})) / \sim_n)$ which belongs to $R_i(\mathbb{L}_n \times \mathbb{A}^2 / \sim_n)$ since κ, α, β and the quotient map from $\mathbb{L}_n \times \mathbb{A}^2$ to $\mathbb{L}_n \times \mathbb{A}^2 / \sim_n$ are homomorphisms. If there exists an index \hat{j} such that $z_{\hat{j}} = x$, then $\kappa(z_j) = 0$ for $j = 1, \dots, r_i$ whence $\phi(z_j) = (0, \alpha(z_j), \alpha(z_j)) / \sim_n$ for $j = 1, \dots, r_i$ by definition of \sim_n ; therefore $(\phi(z_1), \dots, \phi(z_{r_i}))$ is equal to $((0, \alpha(z_1), \alpha(z_1)) / \sim_n, \dots, (0, \alpha(z_{r_i}), \alpha(z_{r_i})) / \sim_n)$ which is in $R_i(\mathbb{L}_n \times \mathbb{A}^2 / \sim_n)$. Similarly if there exists an index \hat{j} such that $z_{\hat{j}} = y$, then $(\phi(z_1), \dots, \phi(z_{r_i})) = ((n, \beta(z_1), \beta(z_1)) / \sim_n, \dots, (n, \beta(z_{r_i}), \beta(z_{r_i})) / \sim_n) \in R_i(\mathbb{L}_n \times \mathbb{A}^2 / \sim_n)$. Thus ϕ is a homomorphism.

Since there exists a homomorphism from an obstruction of \mathbb{A} to $\mathbb{L}_n \times \mathbb{A}^2 / \sim_n$ we conclude that there is no homomorphism from $\mathbb{L}_n \times \mathbb{A}^2 / \sim_n$ to \mathbb{A} . \square

By Theorem 2.5, Corollary 4.3 and Proposition 4.1 we have the following characterisations:

Theorem 4.7. Let \mathbb{A} be a σ -structure. Then the following are equivalent.

1. \mathbb{A} -CSP is first-order definable;
2. For some n there exists a homomorphism from ${}^1\mathbb{A}^n$ to \mathbb{A} ;
3. For some n there exists a homomorphism from $\mathbb{L}_n \times \mathbb{A}^2 / \sim_n$ to \mathbb{A} .

At first glance our situation vis-a-vis the decidability question appears no better than before, but a closer look at the third condition in the above theorem reveals an upper bound on n : indeed, for $0 \leq k \leq n$, the restriction ϕ_k of a homomorphism $\phi : \mathbb{L}_n \times \mathbb{A}^2 / \sim_n \rightarrow \mathbb{A}$ to $\{k\} \times \mathbb{A}^2 / \sim_n$ corresponds to a homomorphism from \mathbb{A}^2 to \mathbb{A} , and there are at most $|A|^{|A|^2}$ of these. If for $k < k'$ we have $\phi_k = \phi_{k'}$, then for $n' = n - k' + k$ we can define a homomorphism $\phi' : \mathbb{L}_{n'} \times \mathbb{A}^2 / \sim_{n'} \rightarrow \mathbb{A}$ by removing the useless middle part. Therefore to determine whether \mathbb{A} -CSP is first-order definable it suffices to search for a homomorphism $\phi : \mathbb{L}_n \times \mathbb{A}^2 / \sim_n \rightarrow \mathbb{A}$ with $n \leq |A|^{|A|^2}$, and this is a finite decision procedure.

We can refine this argument by defining a graph structure on the set of all homomorphisms from \mathbb{A}^2 to \mathbb{A} , where two homomorphisms ψ, ψ' are called *adjacent* if there exists a homomorphism $\phi : \mathbb{L}_1 \times \mathbb{A}^2 \rightarrow \mathbb{A}$ such that $\phi_0 = \psi$ and $\phi_1 = \psi'$. A homomorphism from $\mathbb{L}_n \times \mathbb{A}^2 / \sim_n$ to \mathbb{A} then corresponds to a link of length n between a homomorphism $\phi_0 : \mathbb{A}^2 \rightarrow \mathbb{A}$ which factors through the first projection and a homomorphism $\phi_n : \mathbb{A}^2 \rightarrow \mathbb{A}$ which factors through the second projection. Since undirected reachability can be solved in logarithmic space, in our exponential setting this means that the search can be performed in polynomial space. In the next section this idea is developed further and we prove that the problem of determining whether \mathbb{A} -CSP is first-order definable is actually in **NP**.

5. DISMANTLABILITY

5.1. Preliminaries. Let $\mathbb{A} = \langle A; R_1(\mathbb{A}), \dots, R_m(\mathbb{A}) \rangle$ be a σ -structure. For $x, y \in A$ we say that y *dominates* x in \mathbb{A} , if for every $R_i \in \sigma$, $j \in \{1, \dots, r_i\}$ and $(x_1, \dots, x_{r_i}) \in R_i(\mathbb{A})$ with $x_j = x$ we also have $(y_1, \dots, y_{r_i}) \in R_i(\mathbb{A})$ with $y_j = y$ and $y_k = x_k$ for all $k \neq j$. For instance, if R_i is ternary and $(x, t, x) \in R_i(\mathbb{A})$, then for y to dominate x we must have $(y, t, x) \in R_i(\mathbb{A})$ and $(x, t, y) \in R_i(\mathbb{A})$, each of which also implies $(y, t, y) \in R_i(\mathbb{A})$. We say that x is *dominated* in \mathbb{A} if it is dominated by some element $y \in A \setminus \{x\}$. We say that

\mathbb{A} *dismantles to* its induced substructure \mathbb{B} if there exists a sequence x_1, \dots, x_k of distinct elements of A such that $A \setminus B = \{x_1, \dots, x_k\}$ and for each $1 \leq i \leq k$ the element x_i is dominated in the structure induced by $B \cup \{x_i, \dots, x_k\}$. In other words, the structure \mathbb{B} can be obtained from \mathbb{A} by successively removing dominated elements; the sequence x_1, \dots, x_k is then called a *dismantling sequence*. Note that if \mathbb{A}_x is the substructure of \mathbb{A} induced by $A \setminus \{x\}$, where x is dominated by y in \mathbb{A} , then we can define a retraction $\rho : \mathbb{A} \rightarrow \mathbb{A}_x$ by putting $\rho(x) = y$ and $\rho(z) = z$ for all $z \neq x$. Using composition we then see that if \mathbb{A} dismantles to \mathbb{B} then \mathbb{B} is a retract of \mathbb{A} (the converse does not hold in general). Our first result shows that “dismantling \mathbb{A} to \mathbb{B} ” can be done greedily.

Lemma 5.1. Let \mathbb{A}, \mathbb{B} be σ -structures such that \mathbb{A} dismantles to \mathbb{B} . Then for every dominated element $a \in A \setminus B$ of \mathbb{A} , the substructure \mathbb{A}_a of \mathbb{A} induced by $A \setminus \{a\}$ dismantles to \mathbb{B} .

Proof. Let x_1, \dots, x_k be a dismantling sequence of \mathbb{A} on \mathbb{B} . Note that for some index j we have $x_j = a$. We will show that by removing x_j and perhaps rearranging the sequence we get a dismantling sequence of \mathbb{A}_a on \mathbb{B} . For $i = 1, \dots, k$ let y_i be an element dominating x_i in the substructure \mathbb{A}_i of \mathbb{A} induced by $B \cup \{x_i, \dots, x_k\}$. Note that for some indices i there may be many choices for y_i , and whenever $y_i \neq a$, y_i also dominates x_i in the substructure of \mathbb{A}_a induced by $B \cup \{x_i, \dots, x_k\} \setminus \{x_j\}$. Thus it suffices to show that for all $i \in \{1, \dots, k\}$, we can select y_i other than a .

Let i be the smallest index such that $y_i = a$, and let b be an element dominating a in \mathbb{A} . Note that if $b \notin \{x_1, \dots, x_{i-1}\}$, then b also dominates x_i in the substructure of \mathbb{A}_a induced by $B \cup \{x_i, \dots, x_k\} \setminus \{x_j\}$, hence we can select $y_i = b$ instead. Thus we can assume that $b = x_{i'}$ for some $i' < i$. We then define a finite increasing sequence i_0, i_1, \dots, i_ℓ by putting $i_0 = i'$, and letting i_{p+1} be the index in $\{i_p + 1, \dots, i - 1\}$ such that $y_{i_p} = x_{i_{p+1}}$ if such an index exists. Then x_i is dominated by a in \mathbb{A}_{i_i} , which is dominated by $b = x_{i_0}$ in \mathbb{A} . For $p = 0, \dots, \ell - 1$, x_{i_p} is dominated by $y_{i_p} = x_{i_{p+1}}$ in \mathbb{A}_{i_p} , and x_{i_ℓ} is dominated by $y_{i_\ell} \neq a$ in \mathbb{A}_{i_ℓ} . If $y_{i_\ell} \neq x_i$, then y_{i_ℓ} also dominates x_i in \mathbb{A}_i hence we can select $y_i = y_{i_\ell}$ instead of $y_i = a$. If $y_{i_\ell} = x_i$, then x_i and $a = x_j$ dominate each other in \mathbb{A}_i . In this case, x_1, \dots, x_{i-1} is a dismantling sequence of \mathbb{A}_a on its substructure induced by $B \cup \{x_i, \dots, x_{j-1}\} \cup \{x_{j+1}, \dots, x_k\}$, which is isomorphic to \mathbb{A}_{i+1} via an isomorphism which fixes B , whence \mathbb{A}_a dismantles to \mathbb{B} . \square

5.2. Exponentiation. Let \mathbb{A} and \mathbb{B} be two σ -structures. The \mathbb{A} -th power of \mathbb{B} is the σ -structure

$$\mathbb{B}^{\mathbb{A}} = \langle B^{\mathbb{A}}; R_1(\mathbb{B}^{\mathbb{A}}), \dots, R_m(\mathbb{B}^{\mathbb{A}}) \rangle,$$

where $B^{\mathbb{A}}$ is the set of all maps from A to B , and for $i = 1, \dots, m$ the relation $R_i(\mathbb{B}^{\mathbb{A}})$ consists of all hyperedges (f_1, \dots, f_{r_i}) such that $(f_1(x_1), \dots, f_{r_i}(x_{r_i})) \in R_i(\mathbb{B})$ whenever $(x_1, \dots, x_{r_i}) \in R_i(\mathbb{A})$. This definition is derived from the following correspondence, whose proof is straightforward.

Lemma 5.2. Let $\phi : \mathbb{A} \times \mathbb{C} \rightarrow \mathbb{B}$ be a homomorphism. Then the map $\psi : \mathbb{C} \rightarrow B^{\mathbb{A}}$ defined by $\psi(c) = f_c$, where $f_c(a) = \phi(a, c)$, is a homomorphism from \mathbb{C} to $\mathbb{B}^{\mathbb{A}}$. Conversely, if $\psi : \mathbb{C} \rightarrow \mathbb{B}^{\mathbb{A}}$ is a homomorphism, then the map $\phi : \mathbb{A} \times \mathbb{C} \rightarrow \mathbb{B}$ defined by $\phi(a, c) = \psi(c)(a)$ is a homomorphism from $\mathbb{A} \times \mathbb{C}$ to \mathbb{B} .

In particular the homomorphisms from \mathbb{A} to itself can be viewed as homomorphisms from the product of \mathbb{A} and a loop to \mathbb{A} , which then correspond to loops in $\mathbb{A}^{\mathbb{A}}$.

Now suppose that a is dominated by b in \mathbb{A} , and let ρ be the retraction which maps a to b and fixes every other element of A . Then, considered as an element of $\mathbb{A}^{\mathbb{A}}$, ρ is a “neighbour” of the identity in the sense that there exists a homomorphism ψ from the 1-link \mathbb{L}_1 to $\mathbb{A}^{\mathbb{A}}$ defined by $\psi(0) = \text{id}_A$ and $\psi(1) = \rho$. The main result of this section is a generalisation of this observation to the dismantling process in general.

Lemma 5.3. Let \mathbb{A} be a σ -structure and let \mathbb{B} be a substructure of \mathbb{A} . Then \mathbb{A} dismantles to \mathbb{B} if and only if there exist some $n \geq 0$ and a homomorphism $P : \mathbb{L}_n \rightarrow \mathbb{A}^{\mathbb{A}}$ such that

- (i) $P(0) = \text{id}_A$,
- (ii) B is fixed pointwise by $P(t)$ for every $t = 0, \dots, n$,
- (iii) $P(n)$ is a retraction onto B .

We call two homomorphisms $f, g : \mathbb{A} \rightarrow \mathbb{A}$ *adjacent* if there is a homomorphism P from \mathbb{L}_1 to $\mathbb{A}^{\mathbb{A}}$ such that $P(0) = f$ and $P(1) = g$. Hence Lemma 5.3 states that \mathbb{A} dismantles to \mathbb{B} if and only if there is a link of homomorphisms fixing B pointwise which joins the identity on A to a retraction onto B . The proof will use the following property of composition in powers, whose proof is a straightforward application of the definition.

Lemma 5.4. Let $\mathbb{A}, \mathbb{B}, \mathbb{C}$ be σ -structures. Then the map $\phi : \mathbb{A}^{\mathbb{B}} \times \mathbb{B}^{\mathbb{C}} \rightarrow \mathbb{A}^{\mathbb{C}}$ defined by $\phi(f, g) = f \circ g$ is a homomorphism. In particular for any integer p , the map $\varepsilon_p : \mathbb{A}^{\mathbb{A}} \rightarrow \mathbb{A}^{\mathbb{A}}$ defined by $\varepsilon_p(f) = f^p$ is a homomorphism.

For every $f \in A^A$, and $a \in A$, there exist integers $0 \leq i < j \leq |A|$ such that we have $f^j(a) = f^i(a)$; we say that a has *finite period under f* if we can take $i = 0$. For $p = |A|!$, we then have $f^p(a) = a$ if a has finite period under f , and otherwise $f^p(a)$ has finite period under f . Thus f^p is a set-theoretic retraction of A onto the set of its elements of finite period under f . Therefore for $p = |A|!$, the homomorphism ε_p defined in Lemma 5.4 is a retraction of $\mathbb{A}^{\mathbb{A}}$ onto its substructure induced by the set-theoretic retractions of A .

Proof of Lemma 5.3. Suppose that \mathbb{A} dismantles to \mathbb{B} , and let x_1, \dots, x_k be a dismantling sequence of \mathbb{A} on \mathbb{B} . For $t = 1, \dots, k$, let $y_t \neq x_t$ be an element dominating x_t in the substructure of \mathbb{A} induced by $B \cup \{x_t, \dots, x_k\}$. We define a sequence $\rho_0, \rho_1, \dots, \rho_k$ of retractions inductively by $\rho_0 = \text{id}_A$, $\rho_t(z) = y_t$ if $\rho_{t-1}(z) = x_t$ and $\rho_t(z) = \rho_{t-1}(z)$ otherwise. Let $P : \mathbb{L}_k \rightarrow \mathbb{A}^{\mathbb{A}}$ be defined by $P(t) = \rho_t$. Then $P(0)$ is the identity, B is fixed by each $P(t)$, and $P(k)$ is a retraction onto B . We show that P is a homomorphism.

For $R_i \in \sigma$, let (t_1, \dots, t_{r_i}) be an element of $R_i(\mathbb{L}_k)$. Then there exists an index $t \in \{1, \dots, k\}$ and a subset J of $\{1, \dots, r_i\}$ such that $t_j = t$ if $j \in J$ and $t_j = t-1$ otherwise. We then have $(P(t_1), \dots, P(t_{r_i})) = (f_1, \dots, f_{r_i})$ where $f_j = \rho_t$ if $j \in J$ and $f_j = \rho_{t-1}$ otherwise. For every $(a_1, \dots, a_{r_i}) \in R_i(\mathbb{A})$, we have $(\rho_{t-1}(a_1), \dots, \rho_{t-1}(a_{r_i})) \in R_i(\mathbb{A})$, since ρ_{t-1} is a homomorphism. Now $(f_1(a_1), \dots, f_{r_i}(a_{r_i}))$ coincides with $(\rho_{t-1}(a_1), \dots, \rho_{t-1}(a_{r_i}))$ except for some possible coordinates in J where y_t replaces x_t . Since $\{\rho_{t-1}(a_1), \dots, \rho_{t-1}(a_{r_i})\} \subseteq B \cup \{x_t, \dots, x_k\}$ and y_t dominates x_t in the substructure of \mathbb{A} induced by that subset, we then have $(f_1(a_1), \dots, f_{r_i}(a_{r_i})) \in R_i(\mathbb{A})$. Thus $(f_1, \dots, f_{r_i}) \in R_i(\mathbb{A}^{\mathbb{A}})$. This shows that P is a homomorphism.

Conversely, suppose that $P : \mathbb{L}_n \rightarrow \mathbb{A}^{\mathbb{A}}$ is a homomorphism such that for $\phi_t = P(t)$, $t = 0, \dots, n$ we have $\phi_0 = \text{id}_A$, B is fixed pointwise by each ϕ_t and ϕ_n is a retraction onto B . Put $p = |A|!$. We define three maps as follows.

- (i) $P' : \mathbb{L}_n \rightarrow \mathbb{A}^{\mathbb{A}}$ is defined by $P'(t) = \rho_t := \phi_t^p$. Thus $P' = \varepsilon_p \circ P$, which is a homomorphism by Lemma 5.4.

- (ii) $P'' : \mathbb{L}_n \rightarrow \mathbb{A}^{\mathbb{A}}$, where $P''(t) = \psi_t$ is defined recursively by $\psi_0 = \rho_0$ and $\psi_t = \psi_{t-1} \circ \rho_t$ for $t = 1, \dots, n$. Since ρ_t is idempotent, $\psi_t = \psi_{t-1} \circ \rho_t = \psi_{t-1} \circ \rho_t \circ \rho_t = \psi_t \circ \rho_t$ is adjacent to $\psi_t \circ \rho_{t+1} = \psi_{t+1}$ by Lemma 5.4, whence P'' is a homomorphism.
- (iii) $P''' = \varepsilon_p \circ P'' : \mathbb{L}_n \rightarrow \mathbb{A}^{\mathbb{A}}$ is a homomorphism by Lemma 5.4.

Note that $P'''(0) = P''(0) = P'(0) = P(0) = \text{id}_A$, and since every $P(t)$ fixes B , $P'''(n) = P''(n) = P'(n) = P(n)$ which is a retraction onto B . Also, for $t = 1, \dots, n$, $\hat{\rho}_t := P'''(t)$ is a retraction whose image $\text{im}(\hat{\rho}_t)$ is contained in that of $\hat{\rho}_{t-1}$. We can then show that every $a \in \text{im}(\hat{\rho}_{t-1}) \setminus \text{im}(\hat{\rho}_t)$ is dominated by $\hat{\rho}_t(a)$ in the substructure \mathbb{A}_{t-1} of \mathbb{A} induced by $\text{im}(\hat{\rho}_{t-1})$. Indeed, for $R_i \in \sigma$ and $(a_1, \dots, a_{r_i}) \in R_i(\mathbb{A}_{t-1})$ such that $a_j = a$ for some index j , we have that $(\hat{\rho}_{t-1}(a_1), \dots, \hat{\rho}_{t-1}(a_{j-1}), \hat{\rho}_t(a_j), \hat{\rho}_{t-1}(a_{j+1}), \dots, \hat{\rho}_{t-1}(a_{r_i}))$ is in $R_i(\mathbb{A})$ since $\hat{\rho}_t$ is adjacent to $\hat{\rho}_{t-1}$, whence $\hat{\rho}_t(a)$ dominates a in \mathbb{A}_{t-1} . Therefore \mathbb{A} dismantles to its substructure induced by $\hat{\rho}_n(A) = B$. \square

5.3. $\mathbb{A}^{(\mathbb{A}^2)}$ and $(\mathbb{A}^2)^{(\mathbb{A}^2)}$. Here we interpret Lemma 4.1 in terms of exponential structures. For a σ -structure \mathbb{A} we denote π_1 and π_2 the two projections of \mathbb{A}^2 on \mathbb{A} . The *diagonal* of \mathbb{A}^2 is its substructure $\Delta_{\mathbb{A}^2}$ induced by $\{(a, a) : a \in A\}$.

Lemma 5.5. Let \mathbb{A} be a σ -structure and n an integer. If there exists a homomorphism $P : \mathbb{L}_n \rightarrow \mathbb{A}^{(\mathbb{A}^2)}$ such that $P(0) = \pi_1$ and $P(n) = \pi_2$, then there exists a homomorphism from $\mathbb{L}_n \times \mathbb{A}^2 / \sim_n$ to \mathbb{A} . If \mathbb{A} is a core, the converse also holds.

Lemma 5.6. Let \mathbb{A} be a σ -structure. If \mathbb{A}^2 dismantles to its diagonal, then for some n there exists a homomorphism $P : \mathbb{L}_n \rightarrow \mathbb{A}^{(\mathbb{A}^2)}$ such that $P(0) = \pi_1$ and $P(n) = \pi_2$. If \mathbb{A} is a core, the converse also holds.

Proof of Lemma 5.5. By Lemma 5.2 a homomorphism $P : \mathbb{L}_n \rightarrow \mathbb{A}^{(\mathbb{A}^2)}$ corresponds to the homomorphism $\phi : \mathbb{L}_n \times \mathbb{A}^2 \rightarrow \mathbb{A}$ defined by $\phi(i, a, b) = P(i)(a, b)$. If $P(0) = \pi_1$ and $P(n) = \pi_2$, then ϕ is constant on every \sim_n -equivalence class, hence we can define a homomorphism $\psi : \mathbb{L}_n \times \mathbb{A}^2 / \sim_n \rightarrow \mathbb{A}$ by $\psi((t, a, b) / \sim_n) = \phi(t, a, b)$.

Conversely, any homomorphism $\psi : \mathbb{L}_n \times \mathbb{A}^2 / \sim_n \rightarrow \mathbb{A}$ can be composed with the quotient map $q : \mathbb{L}_n \times \mathbb{A}^2 \rightarrow \mathbb{L}_n \times \mathbb{A}^2 / \sim_n$ to give a homomorphism $q \circ \psi : \mathbb{L}_n \times \mathbb{A}^2 \rightarrow \mathbb{A}$. By Lemma 5.2, $q \circ \psi$ corresponds to a homomorphism $P : \mathbb{L}_n \rightarrow \mathbb{A}^{(\mathbb{A}^2)}$, and by definition of \sim_n there exist homomorphisms ϕ_1, ϕ_2 from \mathbb{A} to itself such that $P(0) = \phi_1 \circ \pi_1$ and $P(n) = \phi_2 \circ \pi_2$. If \mathbb{A} is a core, then by Lemma 4.1, ϕ_1 and ϕ_2 are both the identity, whence $P(0) = \pi_1$ and $P(n) = \pi_2$. \square

Proof of Lemma 5.6. Suppose that \mathbb{A}^2 dismantles to its diagonal $\Delta_{\mathbb{A}^2}$. By Lemma 5.3, for some n there exists a homomorphism $P : \mathbb{L}_n \rightarrow (\mathbb{A}^2)^{(\mathbb{A}^2)}$ such that $P(0)$ is the identity and $P(n)$ is a retraction on $\Delta_{\mathbb{A}^2}$. We can then define a homomorphism $P' : \mathbb{L}_{2n} \rightarrow \mathbb{A}^{(\mathbb{A}^2)}$ by $P'(t) = \pi_1 \circ P(t)$ and $P'(2n - t) = \pi_2 \circ P(t)$ for $t = 0, \dots, n$. Indeed both definitions of $P'(n)$ coincide since $P(n)$ is a retraction on $\Delta_{\mathbb{A}^2}$, and since $P(0)$ is the identity, $P'(0) = \pi_1$ and $P'(2n) = \pi_2$.

Conversely, for every homomorphism $P : \mathbb{L}_n \rightarrow \mathbb{A}^{(\mathbb{A}^2)}$ such that $P(0) = \pi_1$ and $P(n) = \pi_2$, we can define a homomorphism $P' : \mathbb{L}_n \rightarrow (\mathbb{A}^2)^{(\mathbb{A}^2)}$ by $P'(t) = (P(t), \pi_2)$. Then $P'(0) = (\pi_1, \pi_2)$ is the identity and $P'(n) = (\pi_2, \pi_2)$ is a retraction on $\Delta_{\mathbb{A}^2}$. If \mathbb{A} is a core, then since $\Delta_{\mathbb{A}^2}$ is isomorphic to \mathbb{A} via the canonical isomorphism, the restriction of every

$P(t)$ to $\Delta_{\mathbb{A}^2}$ must coincide with π_2 by Lemma 4.1. Hence for $t = 0, \dots, n$, $P'(t)$ fixes $\Delta_{\mathbb{A}^2}$. Therefore \mathbb{A}^2 dismantles to \mathbb{A} by Lemma 5.3. \square

Let \mathbb{A} be a relational structure such that \mathbb{A}^2 dismantles to $\Delta_{\mathbb{A}^2}$. Then by Lemma 5.6, $\mathbb{A}^{(\mathbb{A}^2)}$ contains a link between the two projections, thus for some n there exists a homomorphism from $\mathbb{L}_n \times \mathbb{A}^2 / \sim_n$ to \mathbb{A} by Lemma 5.5. Hence, by Theorem 4.7, \mathbb{A} -CSP is first-order definable. The converse does not hold in general. However, for any retract \mathbb{B} of \mathbb{A} , \mathbb{A} -CSP is equivalent to \mathbb{B} -CSP. In particular, if \mathbb{B} is the core of \mathbb{A} and \mathbb{A} -CSP is first-order definable, then Theorem 4.7, Lemma 5.5 and Lemma 5.6 imply that \mathbb{B}^2 dismantles to $\Delta_{\mathbb{B}^2}$. Therefore we have proved the following:

Theorem 5.7. A relational structure has a first-order definable CSP if and only if it has a retract whose square dismantles to its diagonal.

6. THE COMPLEXITY OF RECOGNISING FIRST-ORDER DEFINABLE CSP'S

Theorem 6.1. The problem of determining whether a relational structure \mathbb{A} has a first-order definable CSP is **NP**-complete.

In fact, we will show the problem to be **NP**-complete even in the restricted case of directed graphs. We contrast this with the following result:

Theorem 6.2. The problem of determining whether a relational structure \mathbb{A} is a core with a first-order definable CSP can be solved in polynomial time.

In particular, Theorem 6.2 implies that deciding whether an input core structure \mathbb{A} has a first-order definable CSP can be done in polynomial time, but our algorithm does not require a certificate that the input is a core.

Proof of Theorem 6.1. Theorem 5.7 shows that the problem is in **NP**. We will show that 3-SAT reduces to the problem of determining whether a given digraph has first-order definable CSP. Let $\mathcal{I} = \bigwedge_{i=0}^{n-1} (L_{i,1} \vee L_{i,2} \vee L_{i,3})$ be an instance of 3-SAT, where each literal is one of the variables x_1, \dots, x_m or its negation, and (without loss of generality) $L_{i,j} \neq L_{i,j'}$ when $j \neq j'$. We construct a digraph H such that \mathcal{I} is satisfiable if and only if H has first-order definable CSP. The vertex-set of H is $\{0, \dots, n-1\} \times \{1, 2, 3\}$, and there is an arc from (i, j) to (i', j') if and only if $i < i'$ and $L_{i,j}$ is not the negation of $L_{i',j'}$.

Thus the map ϕ from H to the transitive tournament T_n on n vertices defined by $\phi(i, j) = i$ is a homomorphism. Furthermore it is not hard to see that for every tree A which admits a homomorphism $\psi : A \mapsto T_n$, there exists a homomorphism $\hat{\psi} : A \mapsto H$ such that $\psi = \phi \circ \hat{\psi}$. Thus the trees that map to H are precisely those which map to T_n . Since T_n has finite duality [NT00], this means that H has first-order definable CSP if and only if T_n is the core of H by Theorem 2.5.

If \mathcal{I} is satisfiable, then selecting for each i an index j_i such that L_{i,j_i} is true yields a homomorphic image $\{(i, j_i) : 1 \leq i \leq n\}$ of T_n in H . Conversely, if $\{(i, j_i) : 1 \leq i \leq n\}$ is a homomorphic image of T_n in H , then we can consistently deem the literals L_{i,j_i} to be true to find a satisfactory truth assignment of \mathcal{I} . Therefore \mathcal{I} is satisfiable if and only if H has first-order definable CSP. \square

Proof of Theorem 6.2. We first test whether \mathbb{A}^2 dismantles to $\Delta_{\mathbb{A}^2}$. According to Lemma 5.1 this step can be performed in polynomial time using the greedy algorithm. If the answer is negative, then either \mathbb{A} is not a core, or it is a core which does not have a first-order definable CSP. In any case, we output “no” and stop. If the answer is positive, then \mathbb{A} does have first-order definable CSP, but it may not be a core. For each pair $a, b \in A, a \neq b$ we form the quotient $\mathbb{A}_{\{a,b\}}$ of \mathbb{A} under the equivalence which identifies a and b . By Theorem 2.5, \mathbb{A} has tree duality, hence the polynomial consistency-check algorithm (see [FV98]) detects whether $\mathbb{A}_{\{a,b\}}$ admits a homomorphism to \mathbb{A} . If such a homomorphism ϕ exists, then \mathbb{A} admits a homomorphism to its proper substructure $\phi(\mathbb{A}_{\{a,b\}})$ hence it is not a core; we then output “no” and stop. If no homomorphism exists from any quotient $\mathbb{A}_{\{a,b\}}$ to \mathbb{A} , then \mathbb{A} is a core. We then output “yes”. \square

7. PRODUCING SOLUTIONS OF FIRST-ORDER DEFINABLE CSP'S

Let \mathbb{A} be a structure such that \mathbb{A}^2 dismantles to its diagonal. Then \mathbb{A} has a first-order definable CSP; furthermore without loss of generality, we can assume that \mathbb{A} is a core, since adding to the type σ a unary relation for each element A preserves the dismantling of \mathbb{A}^2 to its diagonal. Thus, the hyperedge consistency check algorithm is sufficient to determine whether a structure \mathbb{B} admits a homomorphism to \mathbb{A} . It is possible to find an explicit homomorphism from \mathbb{B} to \mathbb{A} in polynomial time using vertex identifications on a trial and error basis. In this section, we provide an alternative algorithm based on dismantlings of $\mathbb{B} \times \mathbb{A}$. We will use the following variation of Lemma 5.1:

Lemma 7.1. *Let \mathbb{B} be a structure which dismantles to two substructures \mathbb{C} and \mathbb{C}' . If neither of \mathbb{C} and \mathbb{C}' have dominated elements, then \mathbb{C} and \mathbb{C}' are isomorphic.*

Proof. Let x_1, \dots, x_n be a dismantling sequence of \mathbb{B} on \mathbb{C}' . For $i = 1, \dots, n$, let \mathbb{B}_i be the substructure of \mathbb{B} induced by $\{x_i, \dots, x_n\} \cup \mathbb{C}'$. We define a sequence $\mathbb{C}_0 = \mathbb{C}, \mathbb{C}_1, \dots, \mathbb{C}_n$ of structures such that $\mathbb{C}_i = \mathbb{C}_{i-1}$ if x_i is not in the universe of \mathbb{C}_{i-1} , and otherwise \mathbb{C}_i is obtained from \mathbb{C}_{i-1} by replacing the element x_i by an element which dominates it in \mathbb{B}_i . By induction we prove that

\mathbb{B} dismantles to \mathbb{C}_i , which is isomorphic to \mathbb{C} .

Indeed, for $i = 0$ this is given, and the induction step clearly works when $\mathbb{C}_i = \mathbb{C}_{i-1}$. Suppose that \mathbb{C}_i is obtained from \mathbb{C}_{i-1} by replacing x_i by y . Note that y is not already in \mathbb{C}_{i-1} since \mathbb{C}_{i-1} is isomorphic to \mathbb{C} which contains no dominated elements. By Lemma 5.1, there exists a dismantling sequence z_1, \dots, z_m of \mathbb{B}_i on \mathbb{C}_{i-1} . By replacing y by x_i in this sequence, we get a dismantling sequence of \mathbb{B}_i on \mathbb{C}_i , whence \mathbb{B} dismantles to \mathbb{C}_i . Moreover, $\mathbb{C}_{i-1} \cup \mathbb{C}_i$ clearly dismantles to both \mathbb{C}_{i-1} and \mathbb{C}_i , whence x_i and y dominate each other in $\mathbb{C}_{i-1} \cup \mathbb{C}_i$. Therefore \mathbb{C}_{i-1} and \mathbb{C}_i are isomorphic.

Thus, \mathbb{C}' contains a substructure isomorphic to \mathbb{C} . By interchanging the roles of \mathbb{C} and \mathbb{C}' , we conclude that \mathbb{C} and \mathbb{C}' are isomorphic. \square

In a product $\mathbb{B} \times \mathbb{A}$, an element (b, a) is said to be *dominated in the second coordinate* if it is dominated by an element of the form (b, a') . We say that $\mathbb{B} \times \mathbb{A}$ *dismantles in the second coordinate* to its substructure \mathbb{C} if \mathbb{C} can be obtained from $\mathbb{B} \times \mathbb{A}$ by successively removing elements that are dominated in the second coordinate. Note that dismantlings of $\mathbb{B} \times \mathbb{A}$ in the second coordinate can be considered as ordinary dismantlings, by adding to the type σ one unary relation $R_b = \{(b, a) : a \in A\}$ for each $b \in B$. Hence the results of Lemma

7.1 apply, and $\mathbb{B} \times \mathbb{A}$ dismantles in the second coordinate to a structure \mathbb{C} with no elements dominated in the second coordinate. Such a structure \mathbb{C} is unique up to isomorphism. For each $b \in B$, there exists at least one $a \in A$ such that $(b, a) \in C$. If for each $b \in B$, there exists exactly one $a = \phi(b) \in A$ such that $(b, a) \in C$, then \mathbb{C} is the *graph* of the function $\phi : B \rightarrow A$. The latter is a homomorphism from \mathbb{B} to \mathbb{A} precisely when \mathbb{C} is isomorphic to \mathbb{B} .

Theorem 7.2. *Let \mathbb{A} be a structure such that \mathbb{A}^2 dismantles to its diagonal. For a structure \mathbb{B} , let \mathbb{C} be a substructure of $\mathbb{B} \times \mathbb{A}$ obtained by dismantling in the second coordinate until no more elements are dominated in the second coordinate. Then \mathbb{B} admits a homomorphism to \mathbb{A} if and only if \mathbb{C} is the graph of a homomorphism from \mathbb{B} to \mathbb{A} .*

Proof. Obviously, if $\mathbb{B} \times \mathbb{A}$ dismantles in the second coordinate to the graph of a homomorphism ψ from \mathbb{B} to \mathbb{A} , then \mathbb{B} admits a homomorphism to \mathbb{A} . The proof of the converse parallels that of Theorem 5.7. As mentioned in the beginning of this section, we can assume that \mathbb{A} is a core. Let \mathbb{B} be a structure which admits a homomorphism to \mathbb{A} .

We first define the structure $(\mathbb{L}_n \times \mathbb{B} \times \mathbb{A} / \sim_n)^*$ as follows: $\mathbb{L}_n \times \mathbb{B} \times \mathbb{A} / \sim_n$ is the quotient of the product $\mathbb{L}_n \times \mathbb{B} \times \mathbb{A}$ under the equivalence \sim_n defined by

$$(k, b, a) \sim_n (k', b', a') \equiv \begin{cases} (k, b, a) = (k', b', a') \\ \text{or } k = k' = 0 \text{ and } a = a' \\ \text{or } k = k' = n \text{ and } b = b'. \end{cases}$$

Note that since \mathbb{B} admits a homomorphism to \mathbb{A} , the fiber $(\{n\} \times \mathbb{B} \times \mathbb{A} / \sim_n)$ is isomorphic to \mathbb{B} , while $(\{0\} \times \mathbb{B} \times \mathbb{A} / \sim_n)$ is not necessarily isomorphic to \mathbb{A} . We complete the structure of $(\mathbb{L}_n \times \mathbb{B} \times \mathbb{A} / \sim_n)^*$ by adding a copy of \mathbb{A} to the fiber $\{0\} \times \mathbb{B} \times \mathbb{A}$: for each $R \in \sigma$ and $(a_1, \dots, a_r) \in R(\mathbb{A})$, we put $((0, a_1, b) / \sim_n, \dots, (0, a_r, b) / \sim_n) \in R((\mathbb{L}_n \times \mathbb{B} \times \mathbb{A} / \sim_n)^*)$.

As in Lemma 4.6, the substructures of $(\mathbb{L}_n \times \mathbb{B} \times \mathbb{A} / \sim_n)^*$ induced by $\{(k, a, b) / \sim_n : k \neq 0\}$ and $\{(k, a, b) / \sim_n : k \neq n\}$ admit natural homomorphisms to \mathbb{A} and \mathbb{B} respectively, whence both of these admit homomorphisms to \mathbb{A} . Thus if n is larger than the diameter of the minimal obstructions of \mathbb{A} , then there exists a homomorphism

$$\alpha : (\mathbb{L}_n \times \mathbb{B} \times \mathbb{A} / \sim_n)^* \rightarrow \mathbb{A}.$$

Note that α corresponds to a link of homomorphisms $\alpha_k \in \mathbb{A}^{\mathbb{B} \times \mathbb{A}}, k = 0, \dots, n$, where $\alpha_0 = \pi_A$ and $\alpha_n = \phi \circ \pi_B$ for some homomorphism $\phi : \mathbb{B} \rightarrow \mathbb{A}$. We use α to define a link of homomorphisms $\beta_k \in (\mathbb{B} \times \mathbb{A})^{\mathbb{B} \times \mathbb{A}}, k = 0, \dots, n$ by

$$\beta_k(b, a) = (b, \alpha((k, b, a) / \sim_n)).$$

Thus, $\beta_0 = \text{id}_{\mathbb{B} \times \mathbb{A}}$, $\beta_k(b, a) = (b, \beta'_k(b, a)), k = 1, \dots, n-1$ and $\beta_n(b, a) = (b, \phi(b))$. There are two desirable properties which would allow us to reach our conclusion: If β_0, \dots, β_n were a link of retractions such that $\beta_0(A \times B) \supseteq \beta_1(A \times B) \supseteq \dots \supseteq \beta_n(A \times B)$, then by Lemma 5.3 we would have that $\mathbb{B} \times \mathbb{A}$ dismantles on $\beta_n(\mathbb{B} \times \mathbb{A})$. However the current link β_0, \dots, β_n may have neither of these properties. Thus we will repeatedly modify our link through the following two procedures:

- (i) If $\gamma_0, \dots, \gamma_n \in (\mathbb{B} \times \mathbb{A})^{\mathbb{B} \times \mathbb{A}}$ is a link with the same properties as β_0, \dots, β_n above, then for $p = |A|!$ the functions $\rho_0, \dots, \rho_n \in (\mathbb{B} \times \mathbb{A})^{\mathbb{B} \times \mathbb{A}}$ defined by $\rho_k = \gamma_k^p$ form a link of retractions, where $\rho_0 = \text{id}_{\mathbb{B} \times \mathbb{A}}$, $\rho_k(b, a) = (b, \rho'_k(b, a)), k = 1, \dots, n-1$ and $\rho_n(b, a) = (b, \rho'_n(b))$ for some $\rho'_n : B \rightarrow A$. However we do not necessarily have $\rho_0(A \times B) \supseteq \rho_1(A \times B) \supseteq \dots \supseteq \rho_n(A \times B)$.

- (ii) If $\rho_0, \dots, \rho_n \in (\mathbb{B} \times \mathbb{A})^{\mathbb{B} \times \mathbb{A}}$ is a link of retractions with the properties given in (i), we define the sequence $\gamma_0, \dots, \gamma_n \in (\mathbb{B} \times \mathbb{A})^{\mathbb{B} \times \mathbb{A}}$ recursively by $\gamma_0 = \rho_0$ and $\gamma_k = \gamma_{k-1} \circ \rho_k, k = 1, \dots, n$. Then we clearly have $\gamma_0(B \times A) \supseteq \gamma_1(B \times A) \supseteq \dots \gamma_n(B \times A)$. Moreover, $\gamma_0 = \text{id}_{\mathbb{B} \times \mathbb{A}}$, $\gamma_k(b, a) = (b, \gamma'_k(b, a)), k = 1, \dots, n-1$ and $\gamma_n(b, a) = (b, \gamma'_n(b))$ for some $\gamma'_n : B \rightarrow A$. In particular, $\gamma_0 = \rho_0 = \text{id}_{\mathbb{B} \times \mathbb{A}}$ and $\gamma_1 = \rho_1$ are adjacent, and if γ_{k-1} and γ_k are adjacent, then so are $\gamma_k = \gamma_{k-1} \circ \rho_k$ and $\gamma_{k+1} = \gamma_k \circ \rho_{k+1}$ by Lemma 5.4. Thus, $\gamma_0, \dots, \gamma_n$ forms a link, though these homomorphisms may not be retractions.

After an initial run through steps (i) and (ii), every time we need to repeat step (ii) it is because the previous step (i) reduced the size of the images of some functions in the link. Thus after some repetitions, we eventually get a link $\delta_0, \dots, \delta_n \in (\mathbb{B} \times \mathbb{A})^{\mathbb{B} \times \mathbb{A}}$ such that $\delta_0 = \text{id}_{\mathbb{B} \times \mathbb{A}}$, each δ_k is a retraction and $\delta_0(A \times B) \supseteq \delta_1(A \times B) \supseteq \dots \supseteq \delta_n(A \times B)$. Thus the map $P : \mathbb{L}_n \rightarrow (\mathbb{B} \times \mathbb{A})^{\mathbb{B} \times \mathbb{A}}$ defined by $P(k) = \delta_k$ satisfies the hypotheses of Lemma 5.3 whence $A \times B$ dismantles to $\delta_n(\mathbb{B} \times \mathbb{A})$. Moreover, each function δ_k is of the form $\delta_k(b, a) = (b, \delta'_k(b, a))$, and also preserve the relations $R_b = \{(b, a) : A \in A\}, b \in B$. Thus $A \times B$ dismantles in the second coordinate to $\delta_n(\mathbb{B} \times \mathbb{A}) = \{(b, \psi(b)) : b \in B\}$ for some function $\psi : B \rightarrow A$. Since there exists a homomorphism $\phi : \mathbb{B} \rightarrow \mathbb{A}$ and the dismantling sequence induces a homomorphism from the graph of ϕ to that of ψ , we conclude that ψ is indeed a homomorphism from \mathbb{B} to \mathbb{A} . \square

Given a structure \mathbb{A} such that \mathbb{A}^2 dismantles to its diagonal, Theorem 7.2 provides the following algorithm for deciding whether a structure \mathbb{B} admits a homomorphism to \mathbb{A} : We dismantle $\mathbb{B} \times \mathbb{A}$ in the second coordinate until we get a structure \mathbb{C} with no dominations in the second coordinate. We then have the following possibilities:

- (i) If \mathbb{C} is not a graph, then there is no homomorphism from \mathbb{B} to \mathbb{A} .
- (ii) If \mathbb{C} is a graph, $\mathbb{C} = \{(b, \phi(b)) : b \in B\}$ where $\phi : B \rightarrow A$ is not a homomorphism from \mathbb{B} to \mathbb{A} , then there is no homomorphism from \mathbb{B} to \mathbb{A} .
- (iii) Otherwise, \mathbb{B} admits a homomorphism to \mathbb{A} , and \mathbb{C} is the graph of such a homomorphism $\phi : \mathbb{B} \rightarrow \mathbb{A}$.

This algorithm works a bit like the hyperedge consistency check, with the list of an element b of B identified with the fiber $\{(b, a) : a \in A\}$. In the dismantling algorithm, an element is removed from a list if it becomes redundant rather than inconsistent. Both algorithms work in $O(|B|^{d+2})$ time, where d is the maximum arity in σ .

8. INFERRED CONSTRAINTS AND L-COMPLETE CSP'S

In this section we analyse the computational complexity of CSP's whose basic relations are inferred from those of a first-order definable CSP. Let Γ be a set of relations on the finite set A . The *relational clone generated by* Γ , denoted by $\langle \Gamma \rangle$, is the set of relations on A inferred from the relations in Γ , i.e. definable from relations in Γ via primitive positive formulas. We now give equivalent combinatorial and algebraic descriptions of the relations in $\langle \Gamma \rangle$ (see e.g. [CJ06]). Recall from Section 4.2 that an operation f on a set A *preserves* a relation θ on A if f is a homomorphism from \mathbb{A}^n to \mathbb{A} where $\mathbb{A} = \langle A; \theta \rangle$.

Lemma 8.1. Let Γ be a finite set of relations on A and let θ be a k -ary relation on A . Then the following conditions are equivalent:

- (1) $\theta \in \langle \Gamma \rangle$;

- (2) every operation on A that preserves every relation in Γ also preserves θ ;
- (3) there exists a (primitive positive) formula

$$\phi(x_1, \dots, x_k) \equiv \exists y_1, \dots, \exists y_m \psi(x_1, \dots, x_k, y_1, \dots, y_m)$$

where ψ is a conjunction of atomic formulas with relations in $\Gamma \cup \{=\}$ such that $(a_1, \dots, a_k) \in \theta$ if and only if $\phi(a_1, \dots, a_k)$ holds;

- (4) there exists a structure \mathbb{X} of the same signature as the structure $\mathbb{A} = \langle A; \Gamma \rangle$, and elements $x_1, \dots, x_k \in X$ such that

$$\theta = \{(f(x_1), \dots, f(x_k)) : f : \mathbb{X} \rightarrow \mathbb{A} \text{ a homomorphism}\}.$$

□

A relation θ of arity $k \geq 2$ is *redundant* if there exist indices $i < j$ such that $x_i = x_j$ for any tuple $\bar{x} \in \theta$; otherwise we say that θ is *irredundant*. If there exist indices $i < j$ such that $x_i = x_j$ for any tuple $\bar{x} \in \theta$, and furthermore there exist at least two distinct values a and b that appear as the i -th coordinate of tuples in θ , then we say that θ is *biredundant*. Stated differently, θ is biredundant if the projection of θ onto two indices yields the equality relation on a set with at least 2 elements.

Theorem 8.2. Let \mathbb{A} be a core structure such that \mathbb{A} -CSP is first-order definable, and let \mathbb{B} be a structure whose basic relations are contained in the relational clone generated by the basic relations of \mathbb{A} . Then

- (1) The problem \mathbb{B} -CSP is in **L**;
- (2) if \mathbb{B} -CSP is not first-order definable, then it is **L**-complete;
- (3) if none of the basic relations of \mathbb{B} is biredundant then \mathbb{B} -CSP is first-order definable; if \mathbb{B} is a core the converse holds as well.

Proof. The first two statements follow from Theorem 5 of [ELT07] and Theorem 3.1 of [LT07]. Indeed, the problem $\neg(\mathbb{B}\text{-CSP})$ is definable in symmetric Datalog, which is enough to ensure that \mathbb{B} -CSP is solvable in logspace. Furthermore, every CSP which is not first-order definable is **L**-hard.

For the third statement we argue as follows: suppose first that no basic relation of \mathbb{B} is biredundant. Since \mathbb{A} is a core with first-order definable CSP, by Corollary 4.3 there exists a map f which is a homomorphism from ${}^1\mathbb{A}^n$ to \mathbb{A} . We shall prove that f is also a homomorphism from ${}^1\mathbb{B}^n$ to \mathbb{B} which will conclude the proof by Corollary 4.3. Let $\theta \in \langle \Gamma \rangle$. If θ is irredundant then in the description of θ in Lemma 8.1 (4) we may choose the elements x_1, \dots, x_k to be distinct. Let f_1, \dots, f_{n-1} be homomorphisms from \mathbb{X} to \mathbb{A} yielding tuples in θ , and let $h : X \rightarrow A$ be *any* map. It is easy to see that the map $p = f(f_1, \dots, f_{n-1}, h)$ is a homomorphism from \mathbb{X} to \mathbb{A} , and hence f is 1-tolerant for θ . In the case where θ is redundant, the argument is almost the same: if for some indices we have $x_i = x_j$, since θ is not biredundant, it follows that the value of f_1, \dots, f_{n-1} at x_i and x_j is a unique value, call it a ; since f is a near-unanimity operation by Lemma 4.4, it follows that the value of p at x_i and x_j is the same and so the tuple produced by p is in θ .

Conversely, suppose that \mathbb{B} is a core and that one of its basic relations is biredundant: we shall show that the structure \mathbb{B}^2 does not dismantle to the diagonal. Indeed, suppose that θ is biredundant and without loss of generality suppose that its projection on the first two coordinates is the equality relation on some subset B of A containing elements 0 and 1. Suppose that we have a dismantling of \mathbb{B}^2 : let $A^2 = X_0, \dots, X_k$ be the successive subsets of A^2 obtained by removal of single elements. We prove by induction that for every

i there exists a tuple of the form $((0, 1), (0, 1), \dots) \in \theta(\mathbb{B}^2)$ with all entries in X_i . This is clear for $i = 0$. Now suppose that there is such a tuple $\bar{x} \in \theta(\mathbb{B}^2)$ with all entries in X_i and that X_{i+1} is obtained from X_i by removal of (c, d) . If (c, d) doesn't appear in \bar{x} then we're done; otherwise by definition of dismantling there exists some element $(c', d') \in X_{i+1}$ that dominates (c, d) and so the tuple obtained from \bar{x} by replacing every occurrence of (c, d) by (c', d') is in $\theta(\mathbb{B}^2)$. It is clear that $(c, d) \neq (0, 1)$ because otherwise the tuple $((0, 1), (c', d'), \dots)$ would be in $\theta(\mathbb{B}^2)$ contrary to the fact that $\theta(\mathbb{B}^2)$ is biredundant. Hence there is a tuple of the desired form with entries in X_{i+1} , showing that no dismantling can end in the diagonal. \square

9. CONCLUSION

We have described a simple polynomial-time algorithm that determines if a finite relational structure is a core with first-order definable CSP (Theorem 6.2), and have proved that deciding FO-definability is **NP**-complete (Theorem 6.1). We have also given various characterisations of FO-definable structures in terms of sets of obstructions (Theorem 2.5), and proved that core structures with finite duality admit a 1-tolerant near-unanimity operation (Corollary 4.3 and Lemma 4.4).

Feder and Vardi's Theorem 2.6 shows that the problem of determining whether an input structure \mathbb{A} has tree duality is decidable. In fact the proof of Theorem 6.1 also implies that this problem is NP-hard, but for the moment it is not known to belong to NP or even to P-space. It would be interesting to have these issues resolved.

In the case of first-order definable CSP's, we now have an algorithm which outputs a yes-no answer to the question as to whether an input structure \mathbb{A} has a first-order definable CSP. Using Lemma 2.3, it is possible to modify it so that in the case where \mathbb{A} -CSP is first-order definable, it outputs a first-order sentence $\Phi_{\mathbb{A}}$ such that \mathbb{B} admits a homomorphism to \mathbb{A} if and only if $\Phi_{\mathbb{A}}$ is true on \mathbb{B} . However the upper bound on the length of $\Phi_{\mathbb{A}}$ involves a tower of exponents. It is not clear whether this is realistic; [NT05] reports cases where the length of $\Phi_{\mathbb{A}}$ can be logarithmic in terms of the size of \mathbb{A} , but there are no examples in the direction of the other extreme.

REFERENCES

- [ABISV05] E. Allender, M. Bauland, N. Immerman, H. Schnoor, and H. Vollmer. The complexity of satisfiability problems: Refining Schaefer's theorem. In *Proc. 30th Math. Found. of Comp. Sci. (MFCS'05)*, pages 71–82, 2005.
- [Ats05] A. Atserias. On Digraph Coloring Problems and Treewidth Duality. *Proceedings of the Twentieth Annual IEEE Symp. on Logic in Computer Science (LICS 2005)*, 106–115, 2005.
- [BKJ00] A. Bulatov, A. Krokhin, P. Jeavons. Constraint satisfaction problems and finite algebras. In *ICALP'00*, volume 1853 of *LNCS*, 272–282, 2000.
- [CJ06] D. Cohen and P. G. Jeavons. *The Complexity of Constraint Languages*, in Handbook of Constraint Programming (F. Rossi, P. van Beek, T. Walsh eds.), chapter 8, Elsevier, 2006, 978 pp.
- [Dal05] V. Dalmau. Linear Datalog and bounded path duality of relational structures. *Logical Methods in Computer Science*, 1(1), 2005.
- [DKV02] V. Dalmau, P. Kolaitis, M. Vardi. Constraint satisfaction, bounded treewidth, and finite-variable logics, in *Constraint Programming '02, (LNCS 2002)* 310–326, 2002.
- [ELT07] L. Egri, B. Larose, and P. Tesson. Symmetric Datalog and constraint satisfaction problems in logspace. Submitted, 2007.

- [FV93] T. Feder, M. Y. Vardi, Monotone monadic SNP and constraint satisfaction, in *Proceedings of the 25rd Annual ACM Symposium on Theory of Computing (STOC)*, San Diego, California, (1993), 612–622.
- [FV98] T. Feder, M. Y. Vardi. The Computational structure of monotone monadic SNP and constraint satisfaction: a study through datalog and group theory. *SIAM Journal of Computing* **28**, (1998), 57-104.
- [Jea98] P. Jeavons. On the algebraic structure of combinatorial problems *Theoret. Comput. Sci.* **200** no. 1-2, 185–204, 1998.
- [LT07] B. Larose and P. Tesson. Universal algebra and hardness results for constraint satisfaction problems. *ICALP 2007*, 267-278.
- [NT00] J. Nešetřil, C. Tardif. Duality theorems for finite structures (characterising gaps and good characterisations). *J. Combin. Theory Ser. B* **80**, 2000, 80–97.
- [NT05] J. Nešetřil, C. Tardif. Short answers to exponentially long questions: extremal aspects of homomorphism duality. *SIAM J. Discrete Math.* **19**, 2005, 914–920
- [Ros05] B. Rossman, Existential Positive Types and Preservation under Homomorphisms, in *Proceedings of the Twentieth Annual IEEE Symp. on Logic in Computer Science, (LICS)*, 467–476, 2005.
- [Sch78] T. J. Schaefer. The complexity of satisfiability problems. In *Proc. 10th ACM STOC*, pages 216–226, 1978.